Contracting the Cosmos: Externalities, Capital Markets, and Coasian Bargaining

Shiming Wu

July 2026

Abstract

The growing energy requirements of artificial intelligence (AI) present significant challenges for terrestrial power grids, prompting the technology sector to consider deploying data centers in low Earth orbit. However, operating in the orbital commons introduces an unpriced spatial externality: the endogenous jump-to-ruin risk associated with orbital congestion, known as the Kessler Syndrome. This paper develops a continuous-time stochastic real options game featuring endogenous asset loss and exogenous technological shocks to examine the capital structure and entry timing for commercial space infrastructure. The analysis highlights a theoretical tension: in environments subject to endogenous jump-to-ruin, standard uncoordinated competitive entry increases physical asset density, thereby elevating the probability of collision and potential market failure. Furthermore, the model shows that traditional debt financing is structurally unviable due to zero collateral recovery upon collision, suggesting a pure equity structure. To address the spatial externality, I model two equilibrium regimes. Under a mandatory insurance regime, global reinsurance syndicates function as decentralized regulatory mechanisms, pricing the spatial externality into a firm's cost of capital. In a self-insured open access regime, physical carrying capacity constraints and endogenous debris removal costs incentivize Coasian bargaining. The findings suggest that the unregulated space market tends to consolidate into a coordinated oligopoly, mitigating the tragedy of the commons through private financial constraints without requiring direct sovereign intervention.

1 Introduction

The recent advancement of artificial intelligence models requires substantial computational resources, which are increasingly bounded by terrestrial thermal limits and energy capacity. Next-generation training clusters require continuous baseload power, placing pressure on terrestrial electrical grids and generating environmental externalities such as carbon emissions and water consumption. To navigate these resource bottlenecks, the aerospace and technology sectors are exploring the deployment of space-based data centers. By locating high-density servers alongside space-based solar arrays, firms can theoretically utilize the vacuum of space as a thermal sink while accessing continuous solar energy. However, deploying this infrastructure requires significant capital expenditure in a shared, stateless environment that is susceptible to irreversible physical congestion.

This transition motivates two interconnected research questions at the intersection of corporate finance and the economics of outer space. First, when large-scale infrastructure projects face an endogenous, unrecoverable physical ruin risk (Kessler Syndrome) that limits collateral recovery and complicates traditional debt financing, what is the optimal corporate capital structure? Associated with this is whether private insurance syndicates can function as an efficient substitute for sovereign state subsidies in underwriting this tail-risk. Second, how can decentralized financial contracts be designed to internalize spatial externalities in an environment that largely lacks sovereign jurisdiction and easily enforceable property rights? Specifically, can capital market requirements—such as equity constraints and reinsurance premiums—encourage cooperative behavior and preserve the orbital commons without relying on a centralized global tax authority?

These questions are increasingly relevant to both theoretical finance and environmental mechanism design. Traditional continuous-time real options models [Dixit and Pindyck, 1994, Grenadier, 2002] generally assume that the underlying physical asset remains intact regardless of market entry dynamics. In the orbital environment, however, unregulated market entry generates an unpriced spatial externality: an endogenous jump-to-ruin process. If standard perfect competition unfolds, firms iteratively exercise their investment options until marginal profit approaches zero. In the context of shared orbit, this behavior maximizes physical density and consequently elevates the probability of a cascading debris event. Standard economic models of open access therefore predict suboptimal outcomes or market failure. A self-enforcing financial architecture is needed to balance competitive technological innovation with sustainable spatial density, especially since centralized Pigouvian taxes are difficult to implement globally.

To investigate this dynamic, I develop a continuous-time stochastic differential game featuring endogenous risk parameters, jump-to-ruin dynamics, and Coasian Nash bargaining. The model evaluates how firms determine the timing of irreversible investment under uncertainty when their actions collectively influence the probability of asset destruction. To ground the continuous-time financial valuation in empirical reality, the stochastic variables are calibrated using high-frequency hourly auction data from cloud computing markets. Furthermore, the structural physical parameters governing orbital friction—such as baseline collision risk and carrying capacity—are estimated using astrodynamics databases and corporate regulatory filings.

The analysis yields several practical results regarding the financing of space infrastructure. First, because physical collision results in zero collateral recovery, debt financing capacity approaches zero, necessitating a pure equity or hybrid capital structure. Second, I find that the extreme risk aversion of equity investors inherently encourages aerospace firms to interface with global reinsurance markets. By pooling risk across the orbital environment, these reinsurance syndicates effectively price the spatial congestion externality into the commercial premiums of operators. Finally, I show that firms can mutually reduce these financial frictions by engaging in decentralized Coasian bargaining to form collaborative traffic consortiums. Even in the absence of mandatory insurance, the empirical simulation suggests that operational drag—such as dodging costs and debris removal taxes—creates an endogenous carrying capacity limit.

The remainder of this paper is organized as follows. Section 2 reviews the related literature spanning real options, space economics, and catastrophe insurance. Section 3 provides the institutional background, including terrestrial energy constraints and orbital congestion risks. Section 4 formalizes the capital structure and insured equity value under jump-to-ruin. Then I solve the strategic real options game to derive the cooperative entry threshold. Section 5 models the self-insured equilibrium, estimating structural orbital frictions and defining the capacity limit. Finally, Section 6 offers concluding remarks.

2 Closely Related Literature

This paper bridges corporate finance and the environmental economics of the astropolitical commons. To provide a comprehensive theoretical framework for financing frontier space infrastructure, my work structurally intersects with four distinct strands of literature.

First, this paper extends the literature on continuous-time real options and competitive equilibrium games. Standard real options models, established by Dixit and Pindyck [1994] and expanded into multi-agent strategic frameworks by Grenadier [2002], traditionally assume that the underlying physical asset remains intact regardless of competitive entry. While extensions of real options theory have incorporated exogenous catastrophic shocks and risk-to-ruin [e.g., Koussis et al., 2007], they typically model these jumps as independent Poisson processes entirely unaffected by the strategic choices of the agents. My paper advances this methodology by embedding an endogenous, state-dependent jump-to-ruin parameter representing physical common-pool degradation (Kessler Syndrome). In my model, the catastrophic failure rate is a function of the aggregate spatial density generated by market entry.

Second, this research contributes to the nascent economics of space debris and orbital carrying capacity. Existing literature formally recognizes the Earth's orbit as a fragile common-pool resource severely threatened by runaway debris growth [Liou and Johnson, 2006, Bongers and Torres, 2023, Rao and Rondina, 2025]. Recent analyses underscore the sheer severity of this orbital externality; an OECD Space Forum [2024] report documents over 100 million debris objects threatening thousands of active satellites, warning of a catastrophic tipping point unless policy intervenes. Similarly, Bongers et al. [2026] extensively survey space-economy market failures, emphasizing that ill-defined property rights and congestion require urgent market-based remedies. Integrated macroeconomic models further quantify these massive stakes: Bongers et al. [2025] estimate that without mitigation, debris could cost over 0.5% of world GDP in the long run, while sensible policies could recover up to 0.6% of GDP. Corroborating this, Nozawa et al. [2023] project a nearly 1.95% decline in global GDP absent intervention, revealing an economic Kessler effect where satellite proliferation eventually undermines the orbit's underlying economic value. Furthermore, game-theoretic work demonstrates that decentralized cleanup is structurally inefficient; Klima et al. [2016] find that selfish agents raise removal costs by approximately 9–11% above a coordinated optimum. To mitigate this spatial externality, prevailing economic solutions have heavily relied on centralized Pigouvian taxation. Most notably, Rao et al. [2020] demonstrate that implementing formal orbital-use fees perfectly internalizes the congestion externality, maximizing social welfare. However, this theoretical optimum inherently requires a sovereign global tax authority—an institution that does not legally or practically exist in the stateless astropolitical frontier. My paper bridges this institutional void. I demonstrate that decentralized capital market constraints can replicate the welfare-maximizing effects of orbital-use fees. By formalizing the Coasian threat of reinsurance repricing, I prove that private financial contracts structurally induce rival firms to form consortiums, endogenously suppressing orbital density without the requirement of state intervention.

Third, to achieve this stateless mechanism design, my model integrates classical corporate finance with the law and economics literature regarding catastrophe insurance as a private regulator. While recent institutional proposals, such as those by Adilov and Alexander [2026], show how regulators can artificially engineer a space market using centralized performance bonds and auctions to internalize debris costs, my approach embeds incentives into decentralized private markets. Foundational theories of corporate insurance demand emphasize risk-pooling, financial synergies, and the mitigation of bankruptcy deadweight costs [Mayers and Smith Jr, 1982, Froot, 2001, Leland, 2007]. Building upon this, recent literature establishes that in environments where public regulation is weak, private catastrophe insurers act as surrogate regulators, enforcing safety standards through strict coverage conditions and risk-based pricing [Yin et al., 2011, He and Faure, 2018]. My paper extends this paradigm into the capitalization of space projects. Because unhedged exposure to absolute physical ruin destroys equity valuation, capital market incentivize aerospace firms into global reinsurance treaties. These syndicates subsequently price the spatial externality directly into the firm's cost of capital, allowing the reinsurance market to act as a shadow social planner that enforces spatial safety through financial friction.

Finally, the astropolitical real options game developed in this paper is motivated by the emerging literature documenting the severe terrestrial energy constraints imposed by artificial intelligence. The exponential scaling of Large Language Models (LLMs) requires unprecedented computational power, triggering a massive, non-cyclical surge in global electricity demand. Recent empirical studies highlight that AI data center power consumption is rapidly approaching parity with the fragile power grid infrastructure [De Vries, 2023, Aljbour et al., 2024]. Quantifying this scale, De Vries-Gao [2026] project that AI workloads alone could emit up to 79.7 million tons of CO_2 by 2025, underlining a terrestrial footprint crisis. This dynamic directly motivates proposals to offload computing off-planet; as Howson [2026] outlines, Orbital Data Centers are emerging as a proposed astropolitical solution to sidestep these massive climate and energy externalities. By formally modeling terrestrial compute pricing as a highly volatile Geometric Brownian Motion governed by grid exhaustion limits, my paper provides the first formal corporate finance framework connecting this terrestrial energy bottleneck to the economic justification for off-world AI infrastructure.

Compared to prior work, this paper's contributions are fourfold. First, I extend continuous-time real options models into a unique setting: firms balance irreversible investments against an endogenous physical collision risk. Unlike standard corporate finance models [Dixit and Pindyck, 1994, Leland, 2007], I incorporate an orbital externality that physically vaporizes the underlying asset, rendering traditional debt unviable. Second, I directly connect financial valuation with orbital ecology. I endogenize how debris density evolves as firms launch, capturing the economic Kessler effect. Third, I introduce a stateless mechanism for preserving orbital safety. Expanding upon the regulatory theories of Adilov and Alexander [2026] and the game-theoretic dilemmas of Klima et al. [2016], I propose a financial incentive (insurance deadweight friction) that internalizes the debris externality. Because this mechanism is embedded directly in equilibrium strategic behavior, it aligns private incentives without relying on sovereign taxation. Finally, I bridge space economics with global energy concerns. Following the astropolitical context of Howson [2026], I model an off-world energy bottleneck: as terrestrial green energy collapses under digital demand [De Vries-Gao, 2026], my framework proposes a financial solution.

3 Institutional Background

To understand the economic motivation for capitalizing a $14 billion space-based data center, it is necessary to examine the intersecting crises of terrestrial energy exhaustion, orbital congestion, and the unique mechanics of space liability insurance.

3.1 Terrestrial AI Energy Crises and Climate Externalities

The advancement of generative artificial intelligence has triggered a massive, non-cyclical surge in global electricity demand. Prior to 2022, traditional cloud data centers accounted for a stable 2% to 3% of total U.S. domestic electricity demand. However, driven by the computationally intensive requirements of training Large Language Models (LLMs), this demand is surging exponentially. AI facilities require gigawatts of uninterrupted baseload power, heavily straining the Earth's electrical grids. The Electric Power Research Institute (EPRI) projects data center demand will explode to between 9% and 17% of total U.S. electricity by 2030 [Aljbour et al., 2024].

This demand shock has exposed a critical timing mismatch in infrastructure. Building high-voltage transmission lines and new power plants requires 10 to 15 years due to environmental permitting, whereas technology companies can deploy gigawatt-scale AI server clusters in under three years. Furthermore, operating these systems generates severe climate externalities. Advanced data centers utilize millions of gallons of water daily for thermal cooling, and projections indicate that AI operations could emit up to 80 million tons of CO_2 annually by 2025 [De Vries-Gao, 2026]. Consequently, technology conglomerates are hitting a physical energy wall on Earth, forcing a consideration to off-world computing solutions [Howson, 2026].

3.2 The Emergence of Orbital AI Megaprojects

Space-based AI data centers are rapidly shifting to engineering reality. Space provides near-continuous solar energy, infinite physical real estate, and a massive vacuum to radiate heat.

To capture this emerging market, launch providers are driving massive consolidation to build vertically integrated space superpowers.

Most notably, on February 2, 2026, SpaceX acquired the artificial intelligence startup xAI in an all-stock transaction valued at $1.25 trillion. By integrating xAI's computational architecture with SpaceX's orbital network, the combined entity aims to deploy gigawatt-scale AI compute in space by 2027. To finance this massive capitalization, SpaceX executed an Initial Public Offering (IPO) in June 2026 under the ticker SPCX, raising $75 billion and achieving a valuation exceeding $2 trillion. Rival firms are aggressively following suit; in June 2026, launch provider Rocket Lab announced an $8 billion acquisition of legacy satellite communications giant Iridium, combining launch capabilities with global L-band spectrum to compete directly with Starlink.

The engineering of an orbital AI data center rests on novel architectural shifts. Rather than using ground-based fans and water, space servers rely on colossal thermal radiators spanning up to 100 meters to disperse heat as infrared radiation. They utilize Optical Inter-satellite Links (laser communication) to handle high-bandwidth AI training data, bypassing slow radio frequencies. To achieve cost efficiency, firms rely on reusable heavy-lift vehicles (like SpaceX's Starship, capable of 150-ton payloads at a fraction of historical costs) and are developing Terra Fabs—massive automated factories designed to mass-produce standardized, radiation-tolerant server satellites [SpaceX, 2026].

3.3 Orbital Congestion and the Kessler Syndrome

However, the deployment of sprawling data megaprojects directly threatens Low Earth Orbit (LEO) with an existential congestion crisis. LEO is a highly fragile, common-pool resource currently crowded by tens of thousands of functioning satellites and over 100 million untracked fragments of space debris [OECD Space Forum, 2024].

The primary physical threat in this environment is the Kessler Syndrome—a theoretical tipping point where the density of objects in LEO is high enough that a single collision generates a cloud of hyper-velocity shrapnel [Liou and Johnson, 2006]. This shrapnel destroys nearby satellites, triggering an unstoppable, self-sustaining chain reaction of secondary impacts. If fully triggered, LEO would become unusable for generations, destroying global communications. Space data centers amplify this risk because their massive thermal radiators and solar arrays present enormous physical cross-sections, and their strict low-latency requirements force them to pack tightly into identical, narrow altitude bands (400 km to 600 km).

Low Earth Orbit (LEO) is critically crowded. Over 43,000 human-made objects larger than 10 cm are currently tracked in orbit, including around 15,000 functioning satellites and a massive amount of high-speed debris. Furthermore, an estimated 1.2 million smaller, untrackable fragments pose severe hazards to spacecraft. The density is particularly extreme in a narrow shell about 550 km up, predominantly driven by SpaceX's Starlink network which alone accounts for roughly 10,000 active satellites.

Figure 1: Orbital Congestion in Low Earth Orbit (LEO)

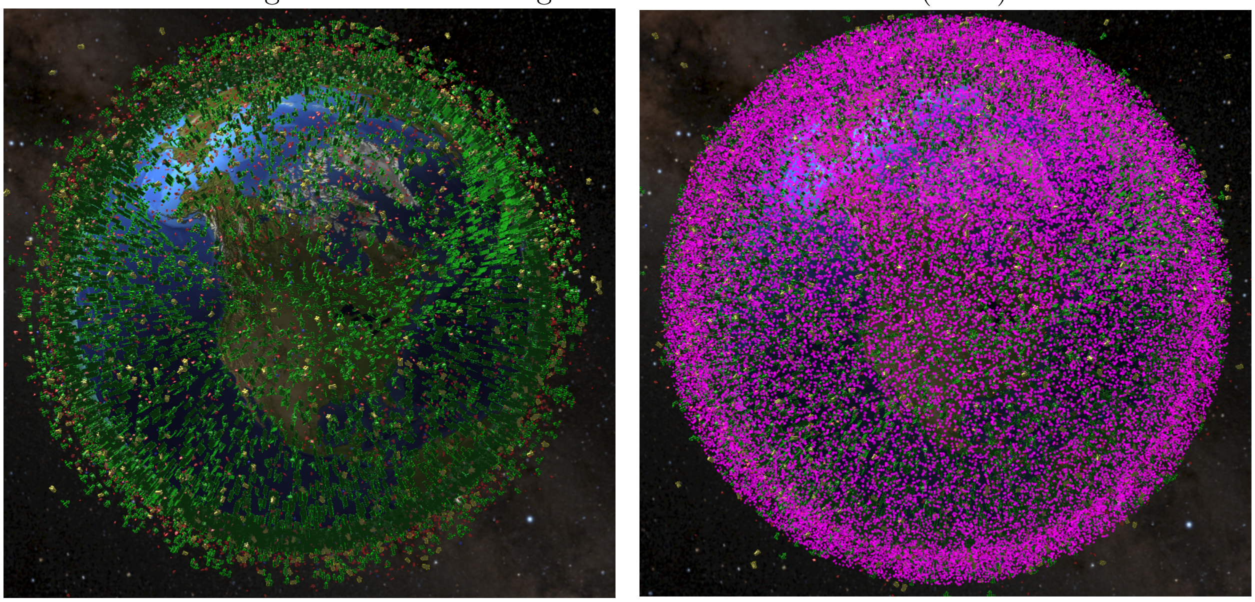

(1) Active payloads and rocket bodies

(2) Including the full debris field

Note: Panel 1 displays actively tracked payloads and rocket bodies. Panel 2 incorporates the full debris field, illustrating the severe spatial capacity constraints and collision risks inherent to the orbital commons. Data source: LeoLabs (July 2026).

Scientists track this urgency using the CRASH clock, which projects how long it would take for a major satellite collision to occur if all operating satellites suddenly lost their ability to steer out of the way. In 2018, this was 164 days; by 2026, it plummeted to just over 2 days. The rapid buildup of inactive spacecraft and fragmented junk moving at up to 28,000 km/h heavily elevates the risk of Kessler Syndrome—a self-sustaining chain reaction of secondary impacts.

Consequently, space physicists utilizing advanced source-sink evolutionary models (such as MOCAT-3) calculate that LEO has a literal carrying capacity. The absolute maximum equilibrium cap for LEO hovers around $\bar{N} = 100,000$ active, coordinated satellites.

3.4 Regulation and Space Insurance

To operate in this hazardous environment, space companies must navigate a complex, two-tiered global insurance framework.

During the launch phase, insurance is strictly mandated and legally enforced. Under 14 CFR Part 440, the U.S. Federal Aviation Administration (FAA) calculates a Maximum Probable Loss (MPL) for third-party liability—often capped around $500 million. Launch providers must purchase this insurance to protect the public from falling debris. If a catastrophic accident exceeds this cap, the U.S. government provides sovereign indemnification.

However, once a satellite achieves orbital insertion, the regulatory landscape shifts dramatically. International space law, such as the UN Liability Convention, holds launching states liable for damages but relies on complex fault-based claims that are difficult to prove in deep space. Consequently, the FAA does not legally mandate continuous in-orbit liability or property insurance. European regulators (e.g., in the UK and Italy) contrast this by moving toward mandatory lifelong in-orbit insurance for their domestic operators.

Despite the lack of a U.S. legal mandate for in-orbit insurance, aerospace companies voluntarily utilize Captive Insurance vehicles and global reinsurance markets to protect their multi-billion-dollar assets. A captive is a wholly owned subsidiary insurer (e.g., SpaceX's Final Frontier Insurance, LLC) that allows the parent company to self-insure routine operational losses using pre-tax cash reserves. However, a catastrophic in-orbit collision could instantly wipe out a captive's liquidity. Therefore, these captives bypass retail insurers and cede the severe tail-risk to the global reinsurance market (e.g., the Lloyd's of London marketplace). By fragmenting the risk into micro-slices across global hedge funds and corporate syndicates, space companies secure the capitalization necessary for launch.

4 Regime I: Reinsurance and Shadow Regulation

In the stateless astropolitical commons, frontier megaprojects generate substantial unpriced spatial externalities. This section argues that, under a mandatory insurance regime, reinsurance pricing can internalize these externalities and mitigate the resulting tragedy of the commons. Because conventional debt financing is largely infeasible in this environment, firms rely heavily on equity financing. The concentration of equity capital increases the demand for effective risk management, elevating the role of global reinsurance markets. By incorporating congestion and collision risks into commercial premiums, reinsurers can function as decentralized social planners, using risk-based pricing to coordinate traffic management and incentivize behaviors that reduce spatial congestion.

4.1 Capital Structure: Pure-Play Startups vs. The Hybrid Company

Consider the astropolitical market where a pioneer (Leader L) deploys space AI data centers. The terrestrial demand for space-based compute, P_t, follows a geometric Brownian motion: $dP_t = \alpha P_t dt + \sigma P_t dW_t$. Firms face a sunk capital expenditure I to launch this infrastructure.

In standard literature, firms optimize their capital structure by balancing debt tax shields against bankruptcy costs. To derive the viability of debt in space, consider a firm attempting to issue a perpetual corporate bond with principal D and continuous coupon c_D. The market risk-free rate is r. The asset is subject to an endogenous collision jump-to-ruin process with intensity λ. Let $R_c \in [0, 1]$ represent the collateral recovery rate upon default.

In continuous-time structural models of corporate debt, the absence of arbitrage implies the existence of a risk-neutral probability measure. Under this measure, the expected instantaneous return of any traded asset equals the risk-free rate r. For a perpetual corporate bond with value B, the return consists of the continuous cash coupon $c_D dt$ and the expected change in the bond's price $\mathbb{E}^Q[dB]$. Because the asset is subject to a Poisson ruin jump (intensity λ) where the bond drops to its collateral recovery value $R_c D$, this no-arbitrage condition is expressed as:

$$rBdt = c_D dt + \lambda(R_c D - B)dt \implies rB = c_D + \lambda(R_c D - B) \tag{1}$$

Case 1: A Pure Space Startup. If an orbital collision occurs, unlike a terrestrial data center where physical servers and real estate can be liquidated during bankruptcy, the collateral recovery rate in space is zero ($R_c = 0$). Substituting $R_c = 0$ yields:

$$rB = c_D - \lambda B \implies B = \frac{c_D}{r + \lambda}$$

Because creditors cannot secure their principal against total collision, the yield to maturity on this debt is $y = c_D/B$. Substituting $B = c_D/(r + \lambda)$, the required credit spread s is:

$$s = y - r = \frac{c_D}{c_D/(r + \lambda)} - r = \lambda \tag{2}$$

Therefore, a spike in orbital congestion ($\lambda \to \infty$) will cause the required credit spread to diverge to infinity, driving the debt capacity of a standalone space project to zero. Consequently, a new space startup structurally relies on a heavy equity capital structure.

Case 2: A Hybrid Company. To relax the zero-debt constraint and finance a multi-billion-dollar frontier AI infrastructure, a firm strategically merges with, or acquires, a mature, highly profitable legacy business (e.g., a satellite internet constellation) that generates a stable and recurring cash flow C.

By integrating these business segments, the combined firm becomes capable of issuing unsecured bonds. Because these bonds are supported by the creditworthiness of the legacy cash flow C, rather than by the physical space-based assets, their repayment is insulated from the Kessler jump risk. As long as the coupon payment satisfies $c_D \leq C$, the firm remains solvent even if the space AI assets are completely destroyed. Consequently, the valuation equation for unsecured debt no longer contains an expected jump-loss term:

$$rBdt = c_D dt + \lambda(B - B)dt \implies B = \frac{c_D}{r}. \tag{3}$$

Therefore, the maximum sustainable debt capacity is simply the present value of the legacy cash flows,

$$D_{\max} = \frac{C}{r}.$$

The remaining multi-billion-dollar capital expenditure (CAPEX) must be financed with equity. Let α denote the expected geometric growth rate (drift) of terrestrial compute demand, governed by the stochastic process

$$dP_t = \alpha P_t dt + \sigma P_t dW_t.$$

Because equity holders bear the full downside risk associated with the space segment, without the protection afforded by debt seniority, they require a higher expected return, denoted by the cost of equity r_E. To ensure standard transversality conditions and rule out explosive asset valuations, I impose the restriction $r_E > \alpha$.

4.2 Endogenous Risk Mitigation and Parametric Captive Insurance

The dependence on equity financing creates a significant pricing friction: unhedged exposure to catastrophic physical tail risk raises the required return on equity, r_E, to prohibitively high levels. To attract private capital, the company must endogenously limit exposure to this jump-to-ruin state. Conventional insurers are generally unwilling to underwrite Kessler-related risks because the associated losses are highly correlated, difficult to quantify, and potentially catastrophic. As an alternative to sovereign subsidies, the company establishes a captive insurance syndicate to internalize and manage this residual tail risk, thereby reducing the required cost of equity and facilitating access to capital markets.

Let λ_i denote the baseline systemic collision intensity when there are i firms operating in the astropolitical market (e.g., λ_1 for a sole operator, λ_2 for a duopoly). Firms can exert continuous Active Debris Removal (ADR) effort $R \geq 0$. Let η represent the marginal efficiency of ADR ($\eta \in (0, \lambda_1/R_{max}]$), yielding a state-dependent residual physical risk of $\lambda_i(R) = \lambda_i - \eta R$. Define the quadratic cost of this effort as $\frac{1}{2}\phi R^2$.

The Captive absorbs standard operational risks, but strictly cedes the unquantifiable Kessler tail-risk to global financial markets via a reinsurance treaty (e.g., with syndicates at Lloyd's of London). If a major collision is detected, the reinsurer pays a fixed indemnification $L = I$. The reinsurance syndicate charges an actuarial premium $\pi_i = \gamma\lambda_i(R)L$, where $\gamma > 1$ represents the reinsurer's capital friction and loading factor.

The firm optimizes its cleanup effort R to minimize the Total Cost function $TC(R)$:

$$\min_R TC(R) = \gamma(\lambda_i - \eta R)L + \frac{1}{2}\phi R^2 \tag{4}$$

Taking the first-order condition with respect to R:

$$\frac{\partial TC}{\partial R} = -\gamma\eta L + \phi R = 0 \implies R^* = \frac{\gamma\eta L}{\phi} \tag{5}$$

Substituting R^* back into the cost function, the minimized baseline cost becomes a constant structural burden M_i:

$$M_i = \gamma\lambda_i L - \frac{1}{2}\frac{(\gamma\eta L)^2}{\phi} \tag{6}$$

To calculate future congestion surcharges, I establish the equity value $E_i(P_t)$. Let A_i represent the firm's compute capacity market share. After servicing the debt coupon c_D, the continuous cash flow to equity holders is $(C - c_D - M_i + A_i P_t)$. The No-Arbitrage condition is:

$$r_E E_i dt = (C - c_D - M_i + A_i P_t)dt + \mathbb{E}[dE_i] \tag{7}$$

To evaluate $\mathbb{E}[dE_i]$, I apply Itô's Lemma for jump-diffusion processes. The underlying terrestrial demand P_t follows the geometric Brownian motion $dP_t = \alpha P_t dt + \sigma P_t dW_t$. Assuming the equity value function $E_i(P_t)$ is twice differentiable and time-homogeneous ($\frac{\partial E_i}{\partial t} = 0$), the continuous diffusion component is derived via a second-order Taylor expansion:

$$\begin{aligned} dE_i^{cont} &= E_i'(P_t)dP_t + \frac{1}{2}E_i''(P_t)(dP_t)^2 \\ &= E_i'(P_t)(\alpha P_t dt + \sigma P_t dW_t) + \frac{1}{2}E_i''(P_t)(\sigma^2 P_t^2 dt) \end{aligned}$$

Taking the expectation, the Wiener process increment vanishes ($\mathbb{E}[dW_t] = 0$), yielding the standard continuous drift: $\mathbb{E}[dE_i^{cont}] = \left(\alpha P_t E_i' + \frac{1}{2}\sigma^2 P_t^2 E_i''\right) dt$.

The asset also faces two independent Poisson jump processes: an endogenous Kessler physical ruin (intensity $\lambda_i(R^*)$) and an exogenous AI algorithmic obsolescence (intensity ξ). Combining the continuous diffusion drift with the expected discrete jumps yields the complete Itô differential:

$$\mathbb{E}[dE_i] = \underbrace{\left(\alpha P_t E_i' + \frac{1}{2}\sigma^2 P_t^2 E_i''\right) dt}_{\text{Continuous GBM Drift}} + \underbrace{\lambda_i(R^*)\left(E_i^{Kessler} - E_i\right) dt}_{\text{Expected Physical Ruin}} + \underbrace{\xi\left(E_i^{AI} - E_i\right) dt}_{\text{Expected AI Obsolescence}} \tag{8}$$

I define the post-jump boundary values. Upon Kessler ruin, space revenues vanish, but the parametric insurance pays L. The debt continues to be serviced by the legacy subsidiary: $E_i^{Kessler} = \frac{C-c_D}{r_E} + L$. Upon an exogenous AI tech shock, space compute becomes obsolete, but because no physical collision occurred, insurance is not triggered: $E_i^{AI} = \frac{C-c_D}{r_E}$.

Substituting the boundary values into the asset pricing condition yields the Hamilton-Jacobi-Bellman (HJB) equation:

$$r_E E_i = (C - c_D - M_i) + A_i P_t + \alpha P_t E_i' + \frac{1}{2}\sigma^2 P_t^2 E_i'' + \lambda_i(R^*)\left(\frac{C - c_D}{r_E} + L - E_i\right) + \xi\left(\frac{C - c_D}{r_E} - E_i\right)$$

I expand the jump terms and move all E_i terms to the left-hand side:

$$(r_E + \lambda_i(R^*) + \xi)E_i = C - c_D - M_i + A_i P_t + \alpha P_t E_i' + \frac{1}{2}\sigma^2 P_t^2 E_i'' + (\lambda_i(R^*) + \xi)\frac{C - c_D}{r_E} + \lambda_i(R^*)L$$

Factoring out $(C - c_D)$ on the right-hand side yields a second-order ODE:

$$(r_E + \lambda_i(R^*) + \xi)E_i = (C - c_D)\left(\frac{r_E + \lambda_i(R^*) + \xi}{r_E}\right) - M_i + \lambda_i(R^*)L + A_i P_t + \alpha P_t E_i' + \frac{1}{2}\sigma^2 P_t^2 E_i''$$

Ruling out speculative financial bubbles, I conjecture a linear fundamental solution $E_i(P) = c_0 + c_1 P$. Taking the derivatives yields $E_i' = c_1$ and $E_i'' = 0$. Substituting these into the ODE allows me to match the P coefficients and the constant terms:

$$P \text{ terms:} \quad (r_E + \lambda_i(R^*) + \xi)c_1 = A_i + \alpha c_1 \implies c_1 = \frac{A_i}{r_E + \lambda_i(R^*) + \xi - \alpha} \tag{9}$$

$$\text{Constants:} \quad (r_E + \lambda_i(R^*) + \xi)c_0 = (C - c_D)\left(\frac{r_E + \lambda_i(R^*) + \xi}{r_E}\right) - M_i + \lambda_i(R^*)L \tag{10}$$

Solving for c_0:

$$c_0 = \frac{C - c_D}{r_E} - \frac{M_i - \lambda_i(R^*)L}{r_E + \lambda_i(R^*) + \xi}$$

Combining c_0 and $c_1 P$ yields the equity value of the insured firm. I define the final term as the insurance friction deadweight (DF_i):

$$E_i(P_t) = \underbrace{\frac{C - c_D}{r_E}}_{\text{Terrestrial Value}} + \underbrace{\frac{A_i P_t}{r_E + \lambda_i(R^*) + \xi - \alpha}}_{\text{Space Compute Value}} - \underbrace{\frac{M_i - \lambda_i(R^*)L}{r_E + \lambda_i(R^*) + \xi}}_{DF_i \text{ (Insurance Deadweight)}} \tag{11}$$

Because M_i contains the reinsurer friction $\gamma > 1$, the deadweight term evaluates to a strict negative, mathematically representing the continuous cost of maintaining the reinsurance.

4.3 Institutional Governance: The Reinsurer as a Decentralized Social Planner

In standard terrestrial markets, a Follower (F) would enter unhindered. In space, uncoordinated entry structurally spikes systemic collision risk to $\lambda_2 \gg \lambda_1$ for all operators. Because the astropolitical commons lacks a sovereign state to levy Pigouvian taxes, I find that global reinsurance portfolios can endogenously act as a decentralized social planner.

First, I establish the Individual Rationality (IR) Constraint for purchasing reinsurance, driven by regulatory capital efficiency. In a regime where in-orbit insurance is mandated, if a firm utilizes a Captive Insurance vehicle, regulators legally require the subsidiary to hold fully funded, liquid cash reserves matching the maximum exposure limit (L). For a space venture funded by high-yield equity, trapping billions of dollars in low-yield reserves incurs a massive continuous opportunity cost of capital equal to $r_E L$.

Alternatively, the captive can purchase an in-orbit reinsurance treaty. The reinsurer charges a continuous premium $\pi_i = \gamma \lambda_i L$, where $\gamma > 1$ represents the reinsurer's capital loading factor and profit margin. The rational firm strictly prefers to cede this risk to the commercial reinsurance market if the continuous reinsurance premium is less than the opportunity cost of trapped equity:

$$(IR): \quad \gamma \lambda_i L < r_E L \implies \gamma \lambda_i < r_E \tag{12}$$

Because the uninsured equity hurdle rate r_E for frontier space ventures vastly exceeds the frictional premium loading ($\gamma \lambda_i$), both the Leader and Follower are mathematically forced to interface with global reinsurance markets.

Second, I establish the Uncoordinated Externality Penalty. Global reinsurers pool risk across the entire astropolitical commons. If the Follower enters the orbit uncoordinated, the systemic collision risk for the orbit spikes to $\lambda_2(\theta, R_F)$. Reinsurers observe this elevated systemic density and reprice the policies for both the Leader and the Follower.

I extract the present value of the Insurance Deadweight Friction (DF) from the fundamental equity valuation (Equation 11). This term represents the continuous frictional cash outflows M_2 (premiums plus debris cleanup) minus expected actuarial payouts $\lambda_2 L$, capitalized at the risk-adjusted discount rate:

$$DF_2(\lambda_2) = \frac{M_2(\lambda_2) - \lambda_2 L}{r_E + \lambda_2 + \xi} \tag{13}$$

By pricing the orbital density risk into the premiums of all operators, the reinsurance syndicate acts as a shadow regulator, translating the physical tragedy of the commons into corporate financial friction.

4.4 Coasian Bargaining and the Open-Data Public Good

Facing punitive reinsurance premiums ($DF_2(\lambda_2)$), both firms are incentivized to lower the systemic risk profile of the orbit. They can form an agreement to coordinate the traffic. If they coordinate, the systemic risk of the orbit drops dramatically to a safe, managed level $\hat{\lambda}_2$.

Let $\Delta E = E^{space}(\hat{\lambda}_2) - E^{space}(\lambda_2) > 0$ denote the gross fundamental enterprise value gained by operating in a safer orbit. Crucially, by coordinating, both firms secure a large reduction in their reinsurance deadweight friction. Let $\Delta DF_{sur} > 0$ denote this mutual actuarial saving:

$$\Delta DF_{sur} = DF_2(\lambda_2) - DF_2(\hat{\lambda}_2) \tag{14}$$

I model this coordination as a Coasian bargain. The Leader built the initial network. To join, the Follower may be asked to pay a licensing fee K. The cooperative Nash surpluses for the Leader (A) and the Follower (B) relative to uncoordinated entry are:

$$\text{Leader's Surplus } (A) = \Delta E + \Delta DF_{sur} + K \tag{15}$$

$$\text{Follower's Surplus } (B) = \Delta E + \Delta DF_{sur} - K \tag{16}$$

Assuming equal bargaining power within the agreement, the firms maximize the Nash Bargaining Product $\Omega(K) = A^{0.5} B^{0.5}$. Taking the log-derivative First-Order Condition with respect to K:

$$\frac{\partial \ln \Omega}{\partial K} = 0.5 \frac{1}{A}\frac{\partial A}{\partial K} + 0.5 \frac{1}{B}\frac{\partial B}{\partial K} = 0 \implies \frac{1}{A}(1) + \frac{1}{B}(-1) = 0 \implies A = B \tag{17}$$

Setting the marginal surpluses equal mathematically isolates the optimal licensing fee:

$$\Delta E + \Delta DF_{sur} + K = \Delta E + \Delta DF_{sur} - K \implies 2K = 0 \implies K = 0 \tag{18}$$

This yields an economic result for astropolitical governance: the optimal side-payment K is zero. Because the risk reduction mutually saves both firms from catastrophic reinsurance repricing, their private financial incentives perfectly align. Capital markets induce competitive firms to reach to an agreement without requiring sovereign intervention or monopolistic toll extraction.

4.5 The Optimal Cooperative Entry Threshold

The Follower optimizes the stopping time τ to enter the market. In the continuation region (where it is optimal to wait), the option to invest $F_2(P)$ yields no cash flows. By the no-arbitrage principle, the expected capital appreciation of this option must equal its risk-adjusted required rate of return. Because an exogenous AI algorithmic obsolescence breakthrough (ξ) would permanently destroy the value of launching the space data center, this Poisson hazard rate is added to the base equity cost (r_E). The no arbitrage condition is:

$$(r_E + \xi)F_2 dt = \mathbb{E}[dF_2] \tag{19}$$

To evaluate $\mathbb{E}[dF_2]$, I apply Itô's Lemma. Given that the underlying terrestrial demand follows $dP = \alpha P dt + \sigma P dW$, I take the second-order Taylor expansion of $F_2(P)$:

$$\begin{aligned} dF_2 &= F_2'(P)dP + \frac{1}{2}F_2''(P)(dP)^2 \\ &= F_2'(P)(\alpha P dt + \sigma P dW) + \frac{1}{2}F_2''(P)(\sigma^2 P^2 dt) \end{aligned}$$

Taking the expectation, the Wiener process increment vanishes ($\mathbb{E}[dW] = 0$). Substituting this back yields the differential equation for the option:

$$(r_E + \xi)F_2 = \alpha P F_2' + \frac{1}{2}\sigma^2 P^2 F_2'' \tag{20}$$

The standard analytical solution is $F_2(P) = kP^\beta$, where k is an endogenous option constant and $\beta > 1$ is the positive root of the quadratic equation: $\frac{1}{2}\sigma^2\beta(\beta - 1) + \alpha\beta - (r_E + \xi) = 0$.

Upon exercising the option, the Follower seamlessly joins the zero-fee agreement. The payoff from exercising is the marginal net present value (NPV) of the new space project. To calculate this, I refer back to the fundamental enterprise valuation in Equation 11, which contains three terms: the terrestrial value, the space compute value, and the insurance friction.

Because the pre-existing legacy terrestrial cash flows ($\frac{C-c_D}{r_E}$) belong to the firm regardless of whether they launch the space megaproject, this constant term strictly cancels out when calculating the marginal added value of the option exercise. The Follower only receives the gross stochastic value added by the new space compute, which is linear in P. I define c_1 as the slope of this fundamental space component:

$$c_1 = \frac{A_2}{r_E + \hat{\lambda}_2 + \xi - \alpha} \implies E_F^{gross\_space}(\hat{\lambda}_2) = c_1 P \tag{21}$$

To maintain algebraic clarity in the real options framework, I isolate the stochastic gross value ($c_1 P$) from all deterministic costs incurred upon entry. I consolidate these costs into a single net effective sunk strike price, I_{eff}. This strike price includes the upfront physical launch cost $I(\theta, h)$—where θ represents the cost-reducing reusable rocketry technology and h represents the target orbital altitude—plus the capitalized Present Value of the Follower's coordinated insurance deadweight friction $DF_2(\hat{\lambda}_2)$ (which acts as the negative constant $-c_0$ from the ODE solution). Summing these yields:

$$I_{eff} = I(\theta, h) + DF_2(\hat{\lambda}_2) = I(\theta, h) + \frac{M_2(\hat{\lambda}_2) - \hat{\lambda}_2 L}{r_E + \hat{\lambda}_2 + \xi} \tag{22}$$

I solve for the critical market entry price threshold $P^*_{F,coop}$ by applying the optimal stopping boundary conditions of continuous-time finance.

First, Value Matching (VM) ensures that at the exact moment of entry $P^*_{F,coop}$, the value of the waiting option equals the net payoff of exercising it.

Second, Smooth Pasting (SP) strictly requires that the marginal derivative of the option value with respect to price ($F_2'(P)$) matches the marginal derivative of the project's net payoff (c_1). In stochastic calculus, if these slopes did not match ($F_2'(P^*) \neq c_1$), the combined value function of the firm would contain a sharp kink at the entry threshold. Under Itô's Lemma, the second derivative at this kink would become a Dirac delta function (approaching infinity). Because the underlying compute demand P_t follows a geometric Brownian motion, it rapidly oscillates back and forth across this threshold, accumulating local time. The infinite second derivative multiplied by the diffusion variance would generate an infinite expected capital appreciation ($\mathbb{E}[dF_2] \to \infty$) over an infinitesimal time step with strictly zero variance.

In financial markets, earning an infinite expected return with zero instantaneous risk constitutes a pure arbitrage opportunity. To strictly prevent this violation of the no-arbitrage theorem, the market forces the slopes to paste smoothly:

$$\text{Value Matching (VM):} \quad k(P^*_{F,coop})^{\beta} = c_1 P^*_{F,coop} - I_{eff} \tag{23}$$

$$\text{Smooth Pasting (SP):} \quad k\beta(P^*_{F,coop})^{\beta-1} = c_1 \tag{24}$$

To analytically solve for $P^*_{F,coop}$, I isolate the option constant k from the Smooth Pasting condition:

$$k = \frac{c_1}{\beta(P^*_{F,coop})^{\beta-1}}$$

Substituting this expression for k into the left side of the Value Matching condition yields:

$$\left(\frac{c_1}{\beta(P^*_{F,coop})^{\beta-1}}\right)(P^*_{F,coop})^{\beta} = c_1 P^*_{F,coop} - I_{eff}$$

Simplifying the exponent on the left-hand side:

$$\frac{c_1 P^*_{F,coop}}{\beta} = c_1 P^*_{F,coop} - I_{eff}$$

I rearrange the terms to isolate the effective sunk cost I_{eff}:

$$I_{eff} = c_1 P^*_{F,coop} - \frac{c_1 P^*_{F,coop}}{\beta} = c_1 P^*_{F,coop}\left(1 - \frac{1}{\beta}\right) = c_1 P^*_{F,coop}\left(\frac{\beta - 1}{\beta}\right)$$

Finally, solving for $P^*_{F,coop}$ provides the exact analytical threshold for market entry:

$$P^*_{F,coop} = \frac{\beta}{\beta - 1}\frac{I_{eff}}{c_1} = \frac{\beta}{\beta - 1}\frac{(r_E + \hat{\lambda}_2 + \xi - \alpha)}{A_2}\left[I(\theta, h) + \frac{M_2(\hat{\lambda}_2) - \hat{\lambda}_2 L}{r_E + \hat{\lambda}_2 + \xi}\right] \tag{25}$$

4.6 Terrestrial Constraints: EIA Data Formulation and Grid Stress

To evaluate the analytical thresholds derived previously and demonstrate the economic magnitude of the spatial externality, I calibrate the model using structural parameters and empirical data. Because space-based AI data centers represent a frontier industry without historical panel data, I employ a structural calibration methodology.

While AI computation pricing (P_t) drives the gross revenue of the space asset, the primary motivation to exercise the real option is driven by the physical exhaustion of the terrestrial grid. To empirically define this constraint, I utilize the U.S. Energy Information Administration (EIA-930) Hourly Electric Grid Monitor.

To isolate the specific impact of hyper-scale AI data centers on the grid, I segment the continuous U.S. electricity data into two distinct cohorts. The High AI Hubs include the Mid-Atlantic (MIDA, home to Data Center Alley in Northern Virginia), Texas (TEX), the Northwest (NW), and California (CAL). These regions possess the highest density of advanced technology infrastructure and hyper-scale server farms. As a control group, I define Low AI Regions comprising New England (NE), New York (NY), Florida (FLA), and Tennessee (TEN), which represent standard residential and commercial load profiles without massive LLM training cluster concentrations.

Let $E_{gen,d}$ represent the total aggregate daily terrestrial power generation capacity across the contiguous United States, and let $E_{dem,d}$ represent the aggregate daily grid demand. The raw daily terrestrial surplus margin is mathematically defined as:

$$S_d = E_{gen,d} - E_{dem,d}$$

Because daily grid operations are subject to extreme high-frequency noise (e.g., localized weather anomalies and weekend industrial dips), I isolate the structural macroeconomic trend by calculating a 30-day moving average of the surplus, denoted as $\bar{S}_d$:

$$\bar{S}_d = \frac{1}{30} \sum_{i=0}^{29} S_{d-i} \tag{26}$$

To systematically identify periods of severe grid vulnerability, I establish a baseline tail-risk metric. Let $\mathcal{D}$ represent the set of all valid operating days in the cleaned 10-year sample, where $N = |\mathcal{D}|$. The 5% critical stress threshold, denoted as $T_{0.05}$, is defined as the 5th empirical percentile of the smoothed surplus margin series $\{\bar{S}_d\}_{d \in \mathcal{D}}$:

$$T_{0.05} = \sup \left\{ s \in \mathbb{R} : \frac{1}{N} \sum_{d \in \mathcal{D}} \mathbf{1}(\bar{S}_d \leq s) \leq 0.05 \right\} \tag{27}$$

where $\mathbf{1}(\cdot)$ is the indicator function. Statistically, $T_{0.05}$ represents a left-tail operating state where the national grid's excess capacity falls into the lowest 5% of its historical distribution.

Figure 2: Regional Electricity Demand Scaling

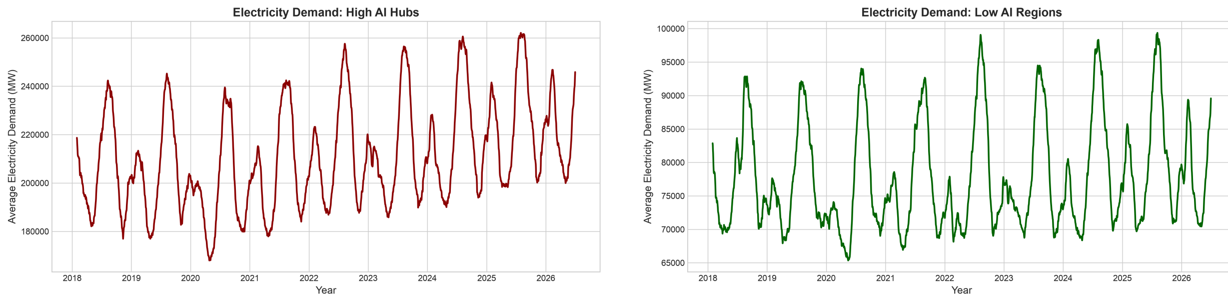

Note: This figure plots the 30-day moving average of daily electricity demand ($\bar{E}_{dem,d}$) across High AI Hubs (left) and Low AI Regions (right). High AI Hubs exhibit a massive, non-cyclical structural drift upward post-2022, driven by hyper-scale AI training loads. Data source: U.S. EIA-930 Hourly Electric Grid Monitor.

The consumption profiles plotted in Figure 2 reveal a divergence between regions saturated with hyper-scale infrastructure and standard grid operations. In the High AI Hubs, average electricity demand exhibits massive baseline scaling, recently cresting 260,000 MW. More importantly, this cohort displays a pronounced structural upward drift that fundamentally accelerates post-2022, mirroring the generative AI boom and the continuous, gigawatt-scale requirements of LLM training workloads. Conversely, the Low AI Regions demonstrate highly seasonal, stable demand patterns. Their peak consumption levels remain firmly anchored below 100,000 MW and lack the aggressive, non-cyclical baseline shift observed in the high-tech hubs, confirming that the demand shock is heavily localized around AI infrastructure.

Figure 3: Terrestrial Grid Surplus Margin and Tail-Risk

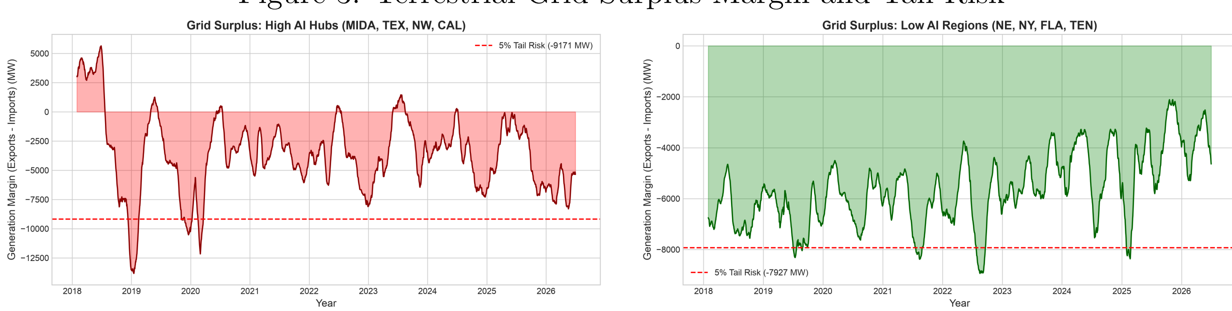

Note: This figure plots the 30-day moving average of the grid generation surplus ($\bar{S}_d$). The dashed line represents the critical 5% historical extreme tail-risk threshold ($T_{0.05}$). The sustained violation of this threshold in High AI Hubs (left panel) highlights the physical terrestrial energy constraints motivating orbital migration.

Analyzing the generation surplus in Figure 3 exposes the severe physical limits capping terrestrial AI expansion. The High AI Hubs are suffering from a steadily deteriorating structural surplus. As these regional grids exhaust their baseline reserves to feed continuous data center draw, the smoothed surplus margin is compressed downward, frequently and persistently violating the historical extreme tail-risk threshold (approximately -9,171 MW) in recent years. This chronic deficit state dictates a heavy, unsustainable reliance on emergency power imports and localized peaker plants. In contrast, the Low AI regions display an improving generation surplus margin over recent years.

The insurmountable timing mismatch between the rapid 3-year deployment of gigawatt-scale AI data centers and the 15-year construction cycles for new power plants means High AI grids cannot generate surplus power fast enough. Because the operating cost of a terrestrial data center is strictly proportional to energy pricing, this localized physical bottleneck forces the frontier AI companies to consider construction of space data center.

4.7 Empirical Estimation

The fundamental stochastic variable in the model, P_t, represents the market-clearing price of frontier AI computation. I calibrate this process using high-frequency hourly auction data from Amazon Web Services (AWS) EC2 Spot Instances. I specifically analyze the `p4d.24xlarge` instance, which provisions 8x NVIDIA A100 GPUs. Because the A100 is the industry-standard semiconductor architecture for training Large Language Models (LLMs), and the AWS spot market utilizes a dynamic auction mechanism, the clearing price serves as a highly efficient, real-time proxy for global AI compute demand. Modeling the asset value P_t as a Geometric Brownian Motion (GBM) perfectly captures the stochastic demand shocks inherent to the AI industry. Because the terrestrial supply of AI supercomputers is relatively fixed in the short-to-medium term, any sudden shifts in the global demand for AI computation immediately translate into price fluctuations.

For the quantitative simulation, I estimate the continuous-time GBM drift (α) and volatility (σ) from the AWS pricing data. However, because AWS spot prices are determined by a dynamic bidding mechanism, raw hourly data exhibits severe high-frequency, mean-reverting microstructure noise. Applying a continuous random walk (GBM) directly to stationary high-frequency noise may artificially inflate the annualized parameters. To correctly isolate the structural macroeconomic trend of the global AI compute transition, I aggregate the hourly auction data into discrete daily average clearing prices.

By Itô's Lemma, the natural logarithm of the computation price $x_t = \ln(P_t)$ follows a generalized Wiener process. Over a discrete time interval Δt, the log-return $X_i = \ln(P_i/P_{i-1})$ is normally distributed. Let $\hat{\mu}$ and $\hat{s}^2$ denote the sample mean and sample variance of the daily log-returns. Because the data is now aggregated daily, the time step is $\Delta t = 1/365$.

Applying Maximum Likelihood Estimation (MLE), the annualized volatility and drift are recovered by scaling the daily variance:

$$\hat{\sigma} = \frac{\hat{s}}{\sqrt{\Delta t}} = \hat{s}\sqrt{365} \tag{28}$$

$$\hat{\alpha} = \frac{\hat{\mu}}{\Delta t} + \frac{1}{2}\hat{\sigma}^2 = \hat{\mu}(365) + \frac{1}{2}\hat{\sigma}^2 \tag{29}$$

To parameterize the astropolitical real options game, I map standard corporate finance metrics and catastrophic insurance multiples to the orbital environment. Table 1 summarizes the baseline calibration with exact economic justifications.

Table 1: Baseline Calibration of Structural Parameters

Parameter	**Value**	**Economic Justification & Source**
Cost of Equity (r_E)	0.15	Aligns with the historical 15.4% average annual return for venture-backed corporate acquisitions [Gompers, 1995].
Tech Obsolescence (ξ)	0.10	Represents a major hardware paradigm shift roughly every 10 years (e.g., optical/quantum computing).
Reinsurance Friction (γ)	1.60	Standard loading factor multiple for catastrophe (Cat) bonds and extreme tail-risk [Froot, 2001].
CAPEX / Limit (I, L)	\$10B	Estimated sunk cost for a heavy-lift Starship payload megaproject [SpaceX, 2026].
Baseline Ruin (λ_2)	0.15	Represents uncoordinated Kessler Syndrome cascade threshold [Liou and Johnson, 2006].
Coordinated Ruin ($\hat{\lambda}_2$)	0.02	Assumed managed residual risk parameter, calibrated to represent successful adherence to ESA space debris mitigation guidelines (e.g., 90% post-mission disposal success) [European Space Agency, 2023].

Note: This table reports the baseline parameters utilized in the continuous-time real options simulation. Parameters map standard venture capital hurdle rates and catastrophe bond loading factors to orbital megaprojects.

Table 1 outlines the structural parameterization driving the investment thresholds. Because space-based infrastructure strictly relies on private equity rather than debt, the cost of equity (r_E) is pegged at 15%, reflecting the historical hurdle rate required by venture-backed aerospace acquisitions. Technological obsolescence ($\xi = 0.10$) acts as an exogenous discount factor, capturing the risk of terrestrial paradigm shifts in AI hardware displacing orbital data centers. Furthermore, I calibrate the baseline physical collision risk ($\lambda_2 = 0.15$) to reflect the uncoordinated Kessler cascade boundary, while managed risk ($\hat{\lambda}_2 = 0.02$) assumes successful adherence to European Space Agency mitigation guidelines. Finally, the insurance deadweight friction ($\gamma = 1.60$) is calibrated to standard commercial loading factors observed in extreme tail-risk catastrophe (Cat) bonds.

By mapping the theoretical continuous-time options game to real-world AI computation pricing and catastrophic reinsurance loading factors, I reveal how capital markets may natively resolve spatial externalities.

Table 2: GBM Parameter Estimation of AI Compute Demand

Estimated GBM Parameter	**Value**
Annualized Drift ($\hat{\alpha}$)	1.5049
Annualized Volatility ($\hat{\sigma}$)	0.7354

Note: This table presents the annualized drift ($\hat{\alpha}$) and volatility ($\hat{\sigma}$) for global AI compute demand. Estimates are derived via Maximum Likelihood Estimation on discrete daily average clearing prices from the AWS EC2 `p4d.24xlarge` spot market, filtering high-frequency microstructure noise.

Table 2 presents the empirical results of the continuous-time demand estimation. By aggregating the data daily, the estimation recovers the macroeconomic parameters. An annualized drift captures the unprecedented, exponential scaling of global demand for AI computation generated by the large language model (LLM) boom. Simultaneously, the volatility ($\sigma \approx 0.735$) aligns with the severe operational uncertainty and pricing oscillations characteristic of frontier hardware spot markets. I feed these empirical parameters into the real options simulation to generate entry thresholds.

Figure 4: Cooperative Entry Threshold vs. Orbital Congestion

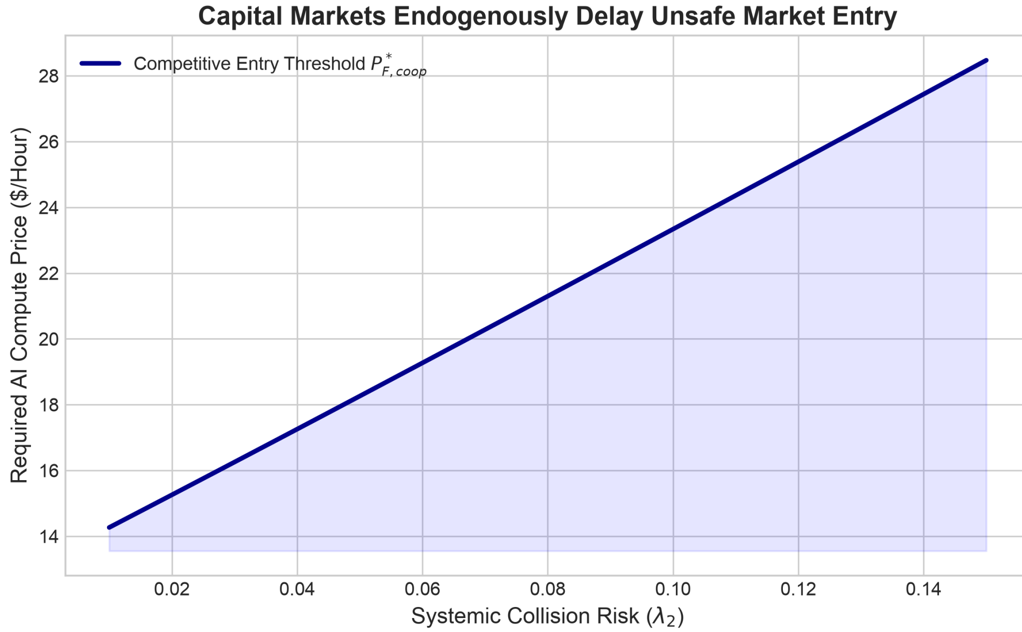

Note: This figure plots the required AI compute clearing price ($P^*_{F,coop}$) to trigger market entry as a function of systemic collision risk (λ_2). The monotonic increase demonstrates that capital markets inherently delay investment and prevent instantaneous density maximization as physical ruin risk escalates.

To bridge the microeconomic hourly spot price (P_t) with the macroeconomic capitalization of the space segment ($I = \$10$ Billion), the model utilizes a scale conversion factor of $A_2 = 0.3066$. Economically, this scales the continuous hourly pricing data into annualized billions of dollars ($35,000$ instances $\times$ $8,760$ hours$/10^9$). This parameter represents a $10 billion orbital megaproject housing the processing capacity of 35,000 frontier high-density server instances, mapping the continuous-time financial valuation into empirically observed terrestrial hardware spot pricing.

As illustrated in Figure 4, uncoordinated perfect competition is disciplined by capital market requirements. The simulation proves that as the systemic physical ruin risk of the orbit (λ_2) increases toward a severe Kessler cascade state, the critical entry threshold required for a follower to exercise their investment option increases monotonically. Under a safely managed baseline risk, firms require approximately $14.20 per hour in computational revenue to justify the astropolitical CAPEX. However, as the spatial externality worsens, the unhedged ruin risk and insurance frictions push the requisite entry price above $28.00 per hour.

Figure 5: Capitalized Insurance Deadweight Penalty

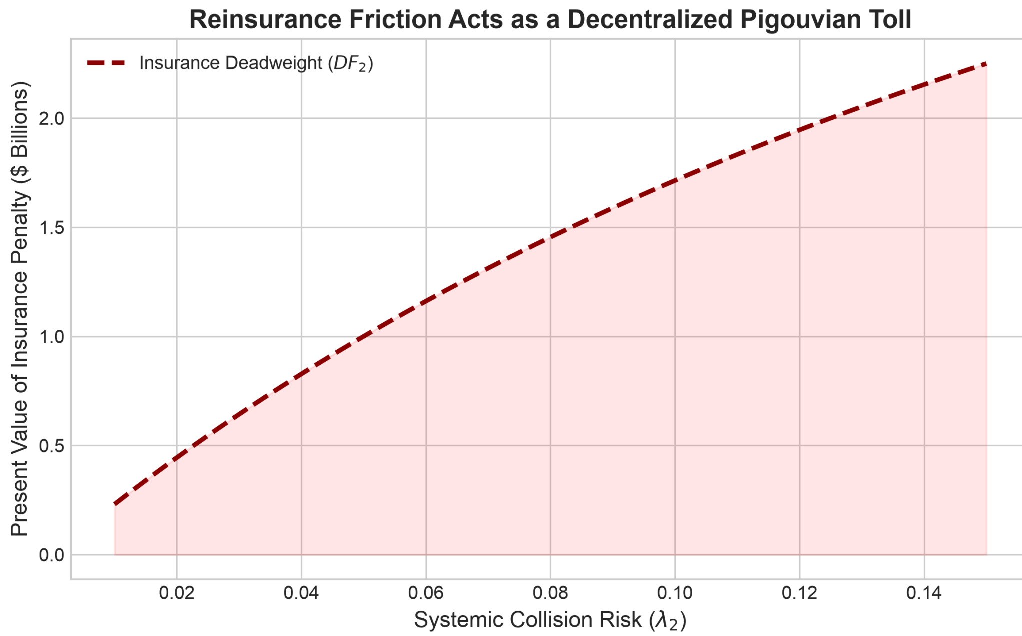

Note: This figure illustrates the present value of the insurance deadweight friction (DF_2) scaled against uncoordinated systemic collision risk (λ_2). The extraction of this penalty acts as a decentralized Pigouvian toll, enforcing traffic coordination through reinsurance pricing.

Figure 5 isolates the fundamental mechanism driving this delay: the decentralized Pigouvian toll extracted by the reinsurance syndicate. Because unhedged absolute ruin destroys equity valuation, capital market forces strictly enforce an Individual Rationality constraint that mandates reinsurance. By accurately pricing the correlated tail-risk of uncoordinated entry into the premiums of all operators, global reinsurers levy a present-value deadweight penalty exceeding $2 Billion under severe congestion scenarios.

Faced with the prospect of mutually assured financial losses arising from reinsurance repricing, competing firms are incentivized to negotiate and coordinate their activities. The credible threat of capitalized financial penalties creates market-based incentives for firms to voluntarily share space domain awareness consortia. This cooperative equilibrium mitigates orbital congestion and supports the off-world energy transition through private-sector coordination rather than direct sovereign intervention.

5 Regime II: The Self-Insured Equilibrium and Coasian Contracts

In Regime I, the mandate for in-orbit insurance internalizes spatial externalities through capitalized reinsurance friction. In practice, however, sovereign governments frequently refrain from mandating in-orbit insurance to incubate domestic aerospace sectors and lower barriers to entry. Consequently, vertically integrated firms attempt to preempt competitors by rapidly deploying satellite networks without global insurance.

Classical Tragedy of the Commons literature suggests that removing this financial penalty triggers runaway entry until the orbit is destroyed. However, in this section, I find that even in the absence of state-mandated insurance, the creation of private Coasian contracts will mitigate congestion. The market organically transitions into a natural oligopoly bounded by an endogenous financial carrying capacity.

5.1 The Tragedy of the Uninsured Commons

Consider an unregulated, self-insured state where N represents the aggregate number of space data centers in a specific Low Earth Orbit (LEO) shell. Firms compete to capture the stochastic terrestrial AI compute demand P_t. Let $R_i \geq 0$ denote the Active Debris Removal (ADR) effort exerted by firm i, with aggregate effort $R_{agg} = \sum_j R_j$. The systemic jump-to-ruin intensity scales exponentially with density but is mitigated by cleanup linearly:

$$\lambda(N, R_{agg}) = \lambda_0 e^{\delta N} - \eta R_{agg} \tag{30}$$

where λ_0 is the baseline risk, $\delta > 0$ is the spatial density multiplier capturing secondary shrapnel generation (Kessler Syndrome), and η is the marginal efficiency of ADR.

Proposition 1 (Open-Access Market Failure): *Under pure, uncoordinated competition, private debris mitigation effort collapses to zero ($R_i^* \to 0$), driving the fundamental equity value of all orbital assets to zero.*

Proof. Firm i operating N_i satellites bears a quadratic private cost of cleanup $\frac{1}{2}\phi R_i^2$. The firm chooses R_i to maximize its fundamental equity value:

$$\max_{R_i} V_i = \frac{N_i P_t}{r_E + (\lambda_0 e^{\delta N} - \eta(R_i + R_{-i})) + \xi - \alpha} - \frac{1}{2}\phi R_i^2 \tag{31}$$

The First-Order Condition (FOC) with respect to R_i yields:

$$\frac{\partial V_i}{\partial R_i} = \frac{N_i P_t \eta}{\left(r_E + \lambda(N, R_{agg}) + \xi - \alpha\right)^2} - \phi R_i = 0 \implies R_i^* = \frac{N_i P_t \eta}{\phi \left(r_E + \lambda(N, R_{agg}) + \xi - \alpha\right)^2} \tag{32}$$

As total market entry $N \to \infty$, the unmitigated baseline risk $\lambda_0 e^{\delta N}$ grows exponentially. Because this exponential term is squared in the denominator of Equation 32, it strictly dominates any linear scaling of the firm's fleet size N_i. Thus, the marginal private benefit of cleanup approaches zero, and $R_i^* \to 0$. Without mitigation, the jump-to-ruin discount rate explodes ($\lambda_u \to \infty$), driving the expected present value of the cash flows to zero: $V_i^{OA} \to 0$. □

5.2 Coasian Bargaining and the Space Traffic Consortium

Because the competitive open access equilibrium creates financial destruction ($V_i^{OA} \to 0$), firms are motivated to bypass the free-rider problem. Without a sovereign state to levy Pigouvian taxes, firms engage in a Coasian bargain, forming a decentralized Space Traffic Consortium.

The Traffic Consortium facilitates the sharing of traffic information among participating operators and jointly finances debris removal activities. Specifically, the Consortium coordinates aggregate commercial active debris removal (ADR) efforts to maintain the orbital collision risk, as measured by the Kessler syndrome hazard rate, at a sustainable target level, $\hat{\lambda}$. The aggregate cleanup effort required to achieve this target is given by

$$R(N) = \frac{\lambda_0 e^{\delta N} - \hat{\lambda}}{\eta},$$

where N denotes the satellite population. Accordingly, the Consortium's total expenditure on debris removal is

$$C_{ADR}(N) = \frac{1}{2}\phi R(N)^2 = \frac{\phi}{2\eta^2}\left(\lambda_0 e^{\delta N} - \hat{\lambda}\right)^2 \tag{33}$$

Proposition 2 (Stability of the Consortium): *The Space Traffic Consortium satisfies both Individual Rationality (IR) and Incentive Compatibility (IC), forming a stable subgame-perfect equilibrium for all $N > N_{critical}$.*

Proof. **(IR Constraint):** A firm joins if the equity value inside the Consortium exceeds the uncoordinated threat point. As established in Proposition 1, uncoordinated entry yields exponential ruin risk, driving the asset pricing discount rate to infinity, thus $V_i^{OA} \to 0$. As

long as terrestrial compute prices P_t yield positive margins under managed risk $\hat{\lambda}$, $V_i^{Coop} > 0$. Thus, $V_i^{Coop} > V_i^{OA}$.
(IC Constraint): A defector avoids paying their share of $C_{ADR}(N)$ but operates blindly, spiking their individual dodging cost to $c^u(N)$ and their collision intensity to $\lambda_u = \lambda_0 e^{\delta N}$.

$$\underbrace{\frac{P_t - c(N) - T_{ADR}(N)}{r_E + \hat{\lambda} + \xi - \alpha}}_{V_i^{Coop}} \geq \underbrace{\frac{P_t - c^u(N)}{r_E + \lambda_0 e^{\delta N} + \xi - \alpha}}_{V_i^{Defect}}$$

For any sufficiently congested orbit ($N > N_{critical}$), the exponential growth of $\lambda_0 e^{\delta N}$ in the defector's discount rate mathematically drives $V_i^{Defect} \to 0$. Defection is strictly dominated.

□

Under asymmetric Nash bargaining, the optimal allocation of total cleanup costs is strictly proportional to fleet size, establishing a uniform per-satellite ADR Tax.

Let $\Phi = \frac{P_t - c(N)}{r_E + \hat{\lambda} + \xi - \alpha}$ represent the capitalized gross per-satellite value. For an Incumbent (I) and Entrant (E), net values are $V_j = \Phi N_j - \frac{C_j}{r}$ where $r = r_E + \hat{\lambda} + \xi$. The asymmetric Nash Bargaining Product, weighted by fleet proportion $\frac{N_j}{N}$, is maximized subject to $C_I + C_E = C_{ADR}(N)$:

$$\max_{C_I} \ln \Omega = \frac{N_I}{N} \ln(V_I) + \frac{N_E}{N} \ln(V_E) \tag{34}$$

Taking the FOC with respect to C_I (noting $\frac{\partial C_E}{\partial C_I} = -1$):

$$\frac{\partial \ln \Omega}{\partial C_I} = \frac{N_I}{N \cdot V_I} \left(\frac{-1}{r} \right) + \frac{N_E}{N \cdot V_E} \left(\frac{1}{r} \right) = 0 \implies \frac{V_I}{N_I} = \frac{V_E}{N_E} \tag{35}$$

Substituting V_j yields $\frac{C_I}{N_I} = \frac{C_E}{N_E}$. Because costs must be shared equally on a per-satellite basis, the negotiated contract equates to a continuous endogenous tax:

$$T_{ADR}(N) = \frac{C_{ADR}(N)}{N} = \frac{\phi}{2N\eta^2} \left(\lambda_0 e^{\delta N} - \hat{\lambda} \right)^2 \tag{36}$$

5.3 Self-Insured Equity Valuation and The Capacity Limit

Because Proposition 2 ensures the Consortium is dynamically stable, we now price the firm's equity strictly under the cooperative regime. As total density N approaches the absolute literal carrying capacity of the orbit ($\bar{N}$), satellites must constantly burn chemical propellant to dodge each other. I model this continuous continuous dodging cost via a standard queuing-theoretic reflecting boundary function: $c(N) = c_0 + \frac{\kappa}{\bar{N} - N}$.

Applying Itô's Lemma, the Hamilton-Jacobi-Bellman (HJB) equation for the self-insured firm incorporates these endogenous operating frictions ($K(N) = c(N) + T_{ADR}(N)$). Solving the resulting second-order ODE yields the fundamental equity value:

$$E_i^{self}(P_t, N) = \underbrace{\frac{C - c_D}{r_E}}_{\text{Terrestrial Value}} + \underbrace{\frac{N_i P_t}{r_E + \hat{\lambda} + \xi - \alpha}}_{\text{Gross Space Compute}} - \underbrace{\frac{N_i \left[c_0 + \frac{\kappa}{\bar{N} - N} + T_{ADR}(N) \right]}{r_E + \hat{\lambda} + \xi}}_{\text{Capitalized STM \& ADR Drag}} \tag{37}$$

When evaluating the real option to launch one additional marginal satellite, the firm faces an effective sunk strike price $I_{eff}^{self}(N)$ that includes the upfront CAPEX $I(\theta, h)$ plus the present value of the lifetime dodging costs and ADR taxes. Applying the standard Value Matching and Smooth Pasting boundary conditions yields the required computational price threshold to trigger market entry:

$$P_{self}^*(N) = \frac{\beta}{\beta - 1} \left(r_E + \hat{\lambda} + \xi - \alpha \right) \left[I(\theta, h) + \frac{c_0 + \frac{\kappa}{\bar{N} - N} + T_{ADR}(N)}{r_E + \hat{\lambda} + \xi} \right] \tag{38}$$

Proposition 3 (The Ultimate Multi-Shell Limit): *Equation 38 establishes an absolute reflecting boundary on capital investment. As $N \to \bar{N}$, the required market entry threshold $P_{self}^*(N) \to \infty$, halting all market entry before the physical carrying capacity is breached.*

Because the ADR tax scales exponentially and the dodging cost scales asymptotically, the marginal expected profit of a new satellite is exhausted by operational drag. Space access organically transitions from an unregulated gold rush into a constrained and low-margin utility market. Under-capitalized startups are mathematically filtered out, leaving a highly coordinated oligopoly that naturally internalizes the spatial externality.

5.4 Estimation and Simulation

To ground the continuous-time financial valuation in empirical reality, I estimate the structural physical parameters $\Theta = \{\kappa, \bar{N}, \lambda_0, \delta\}$ using astrodynamics databases and corporate regulatory filings from 2021 through 2025.

The estimation utilizes three merged data series. First, aggregate orbital density (N_{total}) and the systemic frequency of critical close approaches (x_t) are extracted from the European Space Agency (ESA) Annual Space Environment Reports. Second, to isolate the true operational dodging costs per unit of capital (y_t), I utilize mandatory Semi-Annual Constellation Reports filed by SpaceX to the U.S. Federal Communications Commission (FCC). These reports disclose the number of Collision Avoidance Maneuvers (CAMs) performed by the Starlink fleet. Finally, the active LEO fleet size is cross-verified using the Union of Concerned Scientists (UCS) Satellite Database.

I estimate the structural frictions in two steps. First, the unmitigated spatial density multiplier δ and baseline risk λ_0 are estimated via Poisson Maximum Likelihood. Let x_t be the empirical semi-annual count of critical conjunctions, where $x_t \sim \text{Poisson}(\lambda_0 e^{\delta N_t})$. The log-likelihood function is:

$$\mathcal{L}(\lambda_0, \delta) = \sum_{t=1}^{T} \left(x_t \ln(\lambda_0) + \delta x_t N_t - \lambda_0 e^{\delta N_t} - \ln(x_t!) \right) \tag{39}$$

Second, I estimate the asymptotic dodging friction parameters $\{\kappa, \bar{N}\}$. Let $y_{i,t}$ represent the empirical frequency of CAMs per active satellite for constellation i at time t. Because physical dodging approaches infinity as the orbit fills, the parameters are recovered via Non-Linear Least Squares (NLS) utilizing a standard reflecting boundary function:

$$\{\hat{\kappa}, \hat{\bar{N}}\} = \arg\min_{\kappa, \bar{N}} \sum_{t=1}^{T} \left(y_{i,t} - c_0 - \frac{\kappa}{\bar{N} - N_t} \right)^2 \tag{40}$$

Table 3: Structural Estimation of Orbital Frictions

Parameter	**Estimate**	**Estimation Method**
Panel A: Systemic Risk Frictions		
Baseline Risk (λ_0)	3474.73	Poisson MLE
Density Multiplier (δ)	0.000082	Poisson MLE
Panel B: Operational Dodging Frictions		
Baseline Dodging Cost (c_0)	-51.55	Non-Linear Least Squares
Dodging Sensitivity (κ)	2,677,204.26	Non-Linear Least Squares
Empirical Physical Capacity Limit ($\bar{N}$)	79,216	Non-Linear Least Squares

Note: This table reports the empirical estimates of systemic collision risk and asymptotic dodging frictions. Panel A utilizes Poisson Maximum Likelihood on ESA environmental data. Panel B utilizes Non-Linear Least Squares on SpaceX semi-annual FCC constellation reports (2021-2025).

This estimation strategy provides an empirical bridge. The parameter estimates reported in Table 3 yield consistent alignments with prevailing space physics. The estimated density multiplier ($\delta = 0.000082$) implies that the systemic conjunction risk scales exponentially, multiplying by a factor of e for approximately every 12,000 active satellites introduced to the LEO environment. Furthermore, the Non-Linear Least Squares (NLS) estimation successfully isolates an absolute physical carrying capacity of $\bar{N} \approx 79,216$. This endogenous parameter captures the consensus of astrodynamicists that the orbital shell faces strict, irreversible spatial exhaustion limits below 100,000 objects.

Under the self-insured Incentive Compatible (IC) contract derived in Section 5, the HJB equation incorporates the estimated scaling marginal frictions. Solving the ODE yields the gross space compute value minus the capitalized Space Traffic Management (STM) and ADR drag. Applying the smooth pasting boundary conditions yields the new critical computational price threshold required for market entry:

$$P^*_{self}(N) = \frac{\beta}{\beta - 1} \frac{\left(r_E + \hat{\lambda} + \xi - \alpha\right)}{A_2} \left[I(\theta, h) + \frac{c_0 + \frac{\kappa}{\bar{N} - N} + T_{ADR}(N)}{r_E + \hat{\lambda} + \xi} \right] \tag{41}$$

By plugging the empirical estimates from Table 3 into Equation 41, I simulate the fundamental financial constraints on orbital entry. As shown in Figure 6, the empirical simulation proves the existence of the carrying capacity limit.

Figure 6: The Endogenous Carrying Capacity

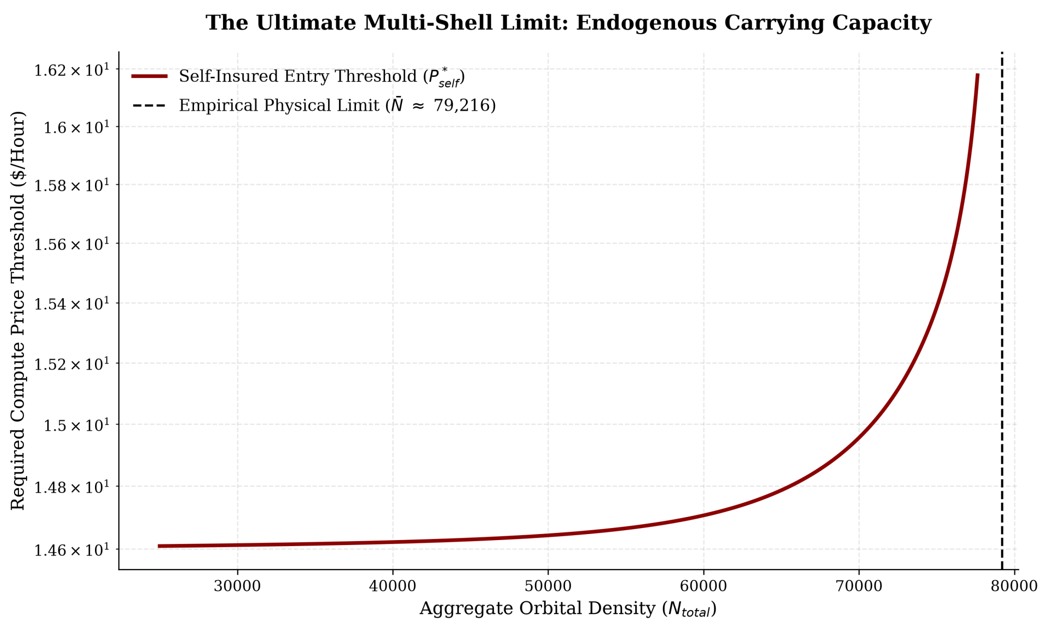

Note: This figure plots the self-insured market entry threshold (P^*_{self}) as aggregate orbital density approaches the empirically estimated limit ($\bar{N} = 79,216$).

While space physicists theorized a maximal physical bound of roughly 100,000 satellites, corporate financial constraints identify a strict empirical limit of $\bar{N} \approx 79,216$. As aggregate density N_{total} scales toward this boundary, the exponential ADR tax (driven by δ) and the asymptotic dodging friction (driven by κ) create an insurmountable reflecting boundary on capital investment. As illustrated in Figure 6, the required entry threshold diverges sharply upward. While a base entry requires approximately \$14.60 per hour in compute pricing, pushing the orbit toward its physical limit forces the required breakeven price rapidly past \$16.20 per hour.

This establishes that even in the total absence of sovereign regulation, the tragedy of the orbital commons is self-correcting through capital market mechanisms. As the orbital shell fills, the marginal expected profit of a new satellite is entirely exhausted by endogenous operational drag. Space access organically transitions from an unregulated open-access gold rush into a heavily constrained, low-margin utility market. Because absorbing these massive continuous friction costs is financially unviable for small operators, the market structurally filters out undercapitalized startups. The resulting astropolitical equilibrium is a natural oligopoly.

6 Conclusion

The potential migration of advanced computing infrastructure to low Earth orbit offers a theoretical pathway to alleviate terrestrial thermal and energy constraints. However, financing these capital-intensive projects requires addressing the fundamental challenge of competition within a fragile, shared environment. This paper develops a continuous-time stochastic real options game with endogenous jump-to-ruin dynamics to examine the optimal financing architecture for commercial space infrastructure.

The analysis indicates that standard uncoordinated market competition in the orbital commons can lead to suboptimal outcomes, as margin-seeking behavior increases physical density and the corresponding probability of collisions. Furthermore, because physical asset loss in space yields no collateral recovery, the model shows that traditional debt financing is structurally unviable, necessitating equity-driven capital structures. To address the underlying spatial externality without relying on a centralized global tax authority, the findings suggest that financial contracts can naturally encourage sustainable practices. Because unhedged exposure to total asset loss diminishes equity valuation, capital market mechanisms incentivize firms to utilize captive insurance vehicles backed by global reinsurance syndicates. These syndicates effectively price the correlated tail-risk of orbital congestion into commercial premiums, functioning as a decentralized regulatory mechanism.

The threat of these capitalized financial frictions provides a strong incentive for rival firms to engage in Coasian bargaining. This dynamic encourages firms to coordinate. Even under a self-insured regime, continuous operational frictions enforce an endogenous carrying capacity, naturally structuring the market into a coordinated oligopoly.

Future research could extend this continuous-time framework to evaluate the financial feasibility of multi-orbit zoning and cislunar infrastructure development. For example, assessing the capital requirements for lunar-based manufacturing and data storage represents a relevant next step in commercial space economics. Additionally, incorporating the frictions of national security interventions and international data-sharing restrictions would provide deeper insight into the institutional constraints shaping the broader space economy.

References

Nodir Adilov and Peter Alexander. Engineering the new space economy: Market creation and institutional design. *Journal of Industrial and Business Economics*, pages 1–19, 2026.

Jordan Aljbour, Tom Wilson, and Poorvi Patel. Powering intelligence: Analyzing artificial intelligence and data center energy consumption. *EPRI White Paper no. 3002028905*, 2024.

Anelí Bongers and José L Torres. Orbital debris and the market for satellites. *Ecological Economics*, 209:107831, 2023.

Anelí Bongers, César Ortiz, and José L Torres. Dise: A dynamic integrated space-economy model for orbital debris mitigation policy evaluation. *Environmental and Resource Economics*, 88(8):2125–2156, 2025.

Anelí Bongers, Benedetto Molinari, Sébastien Rouillon, and José L Torres. The economics of outer space: Regulation, market failures, and industry prospects. *Space Policy*, page 101765, 2026.

Alex De Vries. The growing energy footprint of artificial intelligence. *Joule*, 7(10):2191–2194, 2023.

Alex De Vries-Gao. The carbon and water footprints of data centers and what this could mean for artificial intelligence. *Patterns*, 7(1), 2026.

Avinash K Dixit and Robert S Pindyck. *Investment Under Uncertainty*. Princeton University Press, 1994.

European Space Agency. Esa space environment report 2023. Technical report, ESA Space Debris Office, 2023.

Kenneth A Froot. The market for catastrophe risk: a clinical examination. *Journal of Financial Economics*, 60(2-3):529–571, 2001.

Paul A Gompers. Optimal investment, monitoring, and the staging of venture capital. *The Journal of Finance*, 50(5):1461–1489, 1995.

Steven R Grenadier. Option exercise games: An application to the equilibrium investment strategies of firms. *The Review of Financial Studies*, 15(3):691–721, 2002.

Qihao He and Michael Faure. Regulation by catastrophe insurance: A comparative study. *Connecticut Insurance Law Journal*, 24(1):189, 2018.

Peter Howson. Extra terra nullius: Off-worlding the externalities of ai, bitcoin mining and cloud computing with orbital data centres. *Energy Research & Social Science*, 136:104725, 2026.

Richard Klima, D Bloembergen, R Savani, K Tuyls, D Hennnes, and D Izzo. Game theoretic analysis of the space debris removal dilemma. *ESA report 15*, 8401:2016, 2016.

Nicos Koussis, Spiros H Martzoukos, and Lenos Trigeorgis. Real options with random controls, rare events, and risk-to-ruin. In *Optimisation, Econometric and Financial Analysis*, pages 251–271. Springer, 2007.

Hayne E Leland. Financial synergies and the optimal scope of the firm: Implications for mergers, spinoffs, and structured finance. *The Journal of Finance*, 62(2):765–807, 2007.

J-C Liou and Nicholas L Johnson. Risks in space from orbiting debris. *Science*, 311(5759): 340–341, 2006.

David Mayers and Clifford W Smith Jr. On the corporate demand for insurance. In *Foundations of insurance economics: Readings in economics and finance*, pages 190–205. Springer, 1982.

Wataru Nozawa, Kenichi Kurita, Tetsuya Tamaki, and Shunsuke Managi. To what extent will space debris impact the economy? *Space Policy*, 66:101580, 2023.

OECD Space Forum. *The Economics of Space Sustainability*. OECD Publishing, 2024. URL https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/06/the-economics-of-space-sustainability_5236a39b/b2257346-en.pdf.

Akhil Rao and Giacomo Rondina. The economics of orbit use: Open access, external costs, and runaway debris growth. *Journal of the Association of Environmental and Resource Economists*, 12(2):353–388, 2025.

Akhil Rao, Matthew G Burgess, and Daniel Kaffine. Orbital-use fees could more than quadruple the value of the space industry. *Proceedings of the National Academy of Sciences*, 117(23):12756–12762, 2020.

SpaceX. Starship capability and payload to orbit estimates. https://www.spacex.com/vehicles/starship/, 2026.

Haitao Yin, Howard Kunreuther, and Matthew W White. Risk-based pricing and risk-reducing effort: Does the private insurance market reduce environmental accidents? *The Journal of Law and Economics*, 54(2):325–363, 2011.