

Earth-Orbit Congestion in Constellation Valuation

A Cash Flow Framework for Collision Avoidance Manoeuvre Capacity Insurance Replacement and Terminal Value

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Earth orbit is an operating environment and a finite shared resource. A satellite constellation earns cash through service availability while its spacecraft move through altitude shells that contain active satellites, spent hardware and debris. Congestion changes the economics through conjunction screening, analyst workload, manoeuvres, propellant consumption, interrupted service, shortened asset life, insurance exclusions, replacement demand and regulatory obligations. A valuation that treats collision risk as a remote legal contingency can overstate capacity, free cash flow and terminal value.

This paper develops a congestion-adjusted framework for constellation valuation. It links the tracked environment, operator ephemerides, conjunction data messages, screening thresholds, manoeuvre rules, spacecraft autonomy, propulsion margins, service commitments, insurance and disposal performance to cash flow. The framework separates recurring operational burden from low-frequency losses and from strategic options such as autonomous coordination, improved tracking and lower-risk orbital shells. It also distinguishes operator-controlled risk from hazards created by non-maneuvrable objects and other operators.

Current primary evidence establishes the operating context. ESA reports more than 40,000 tracked objects, about 11,000 active payloads and an estimated population exceeding 1.2 million debris objects larger than one centimetre in its 2025 report. Its 2026 report records more than 300 launches and over 4,000 payloads added during 2025. NASA requires conjunction assessment and collision-avoidance planning across the mission lifecycle. FCC authorisations impose debris, manoeuvrability, disposal and reporting conditions. Public-company filings describe collision, insurance, interruption and replacement exposures. These observations support the framework; they do not determine the value of any unidentified operator. [1][2][3][4][5][6]

The worked case is wholly hypothetical and describes no identified company. It assumes a 4,200-satellite LEO communications constellation, 68,000 screened conjunction messages a year, 2,400 escalated events and 410 avoidance manoeuvres. A USD 8.20 billion headline enterprise value falls to USD 5.05 billion after recognising recurring congestion operations, capacity interruption, shortened life, uninsured collision exposure, replacement liquidity and a lower terminal multiple. A USD 0.55 billion option for automated coordination and improved tracking is admitted only after performance and regulatory milestones; USD 0.20 billion of implementation cost is deducted. Every amount, probability and operating assumption must be replaced with transaction-specific evidence.

The framework is intended for boards, investors, lenders and transaction teams evaluating acquisitions, minority investments, financings and strategic partnerships. Its central output is a bridge from orbital conditions and operating controls to customer capacity, sustainable free cash flow and congestion-adjusted value.

JEL Classification: G12, G31, G32, L96, O32, Q56

Keywords: orbital congestion valuation, constellation economics, collision avoidance, conjunction assessment, space debris, satellite insurance, manoeuvre cost, replacement capital, terminal value, space traffic coordination

Introduction

Constellation valuation usually begins with satellites, capacity, customers and spectrum. Orbit congestion adds a shared environmental constraint that can influence each of them. A close approach consumes analyst attention, computing, coordination and sometimes propellant. A manoeuvre can interrupt payload operations or alter coverage. An impact or fragmentation can remove capacity, create service credits, trigger replacement capital and make a shell more difficult to use.

The relevant question is economic: how much cash flow can the constellation deliver after the recurring cost, capacity loss and capital consequences of operating safely in its orbital environment? That question requires evidence from the operator's conjunction history, flight dynamics, propulsion, insurance, service records and regulatory obligations.

This paper treats congestion as a measurable operating and valuation variable. The method preserves separate ledgers for recurring operations, expected loss, tail liquidity and strategic improvements so that one risk is not counted repeatedly.

1. Define the valuation perimeter

The analysis begins with the legal entities, spacecraft, spectrum, ground systems, launch arrangements, service commitments and disposal obligations that create or absorb congestion exposure. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises an entity-to-asset map, active fleet register, mission licences, customer service levels and third-party operating agreements. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into which cash flows and liabilities travel with the transaction. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should approve a perimeter that reconciles legal ownership, operational control and economic exposure.

2. Reconstruct the orbital environment by shell

Congestion becomes decision-useful when the analyst can trace the altitude, inclination and phasing of the fleet relative to active spacecraft, rocket bodies and debris populations. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain operator ephemerides, public catalogues, regulator filings and independent tracking, reconciled by date and object identifier. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures traffic density and crossing exposure by shell rather than one fleet-wide average. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should identify shells whose congestion can change deployment, replacement or terminal assumptions; the assumption register should show every resulting change in cash, capital and value.

3. Build a conjunction-event ledger

The practical control question concerns the full progression from screened approach through alert, escalation, coordination, manoeuvre decision and post-event review. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes time-stamped conjunction messages, probability updates, covariance quality, counterpart communications and final disposition. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate event volumes, analyst hours, computing load and decision latency. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should separate routine screening from events capable of changing cost, capacity or risk, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

4. Test tracking and covariance quality

The analysis begins with uncertainty in the estimated positions and velocities that drives false alarms, late alerts and decision quality. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises sensor provenance, covariance realism tests, ephemeris age, residual analysis and operator calibration records. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into the effect of data quality on alert volume and unnecessary manoeuvres. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should fund independent validation where weak tracking evidence can change value.

5. Define manoeuvre decision rules

Congestion becomes decision-useful when the analyst can trace the probability, miss distance, consequence and coordination thresholds that determine when the fleet moves. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain approved flight rules, authority matrices, automated recommendations, human approvals and exception logs. Each record needs a stable identifier and a reconciliation to the relevant

spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures manoeuvre frequency and accepted residual risk under consistent rules. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should confirm that thresholds are documented, governed and compatible with licences and service commitments; the assumption register should show every resulting change in cash, capital and value.

6. Cost the conjunction-operations function

The practical control question concerns the people, software, sensors, cloud, communications and assurance required to screen and resolve close approaches. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes headcount, vendor contracts, software licences, system availability, on-call records and event workload. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate a recurring cash cost that scales with fleet, environment and automation quality. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should include a sustainable staffing and systems budget in base free cash flow, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

7. Measure propellant consumption

The analysis begins with the fuel or stored momentum consumed by avoidance, return-to-slot and coordination manoeuvres. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises command histories, achieved delta-v, propulsion efficiency, tank estimates and design reserves. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into life reduction and replacement acceleration by cohort. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should translate manoeuvre policy into remaining life and capital rather than a generic risk premium.

8. Measure service interruption and capacity loss

Congestion becomes decision-useful when the analyst can trace payload outages, beam reallocation, coverage gaps and degraded redundancy before, during and after manoeuvres. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain network telemetry, customer tickets, service credits, traffic rerouting and acceptance records. Each record needs a stable identifier and a reconciliation to the relevant spacecraft,

shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures lost contribution and customer consequences by geography and service class. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should price interruption using contract economics and observed restoration performance; the assumption register should show every resulting change in cash, capital and value.

9. Distinguish avoidable and unavoidable exposure

The practical control question concerns risk created by manoeuvrable neighbours, non-responsive operators, derelict objects and untracked debris. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes coordination records, object manoeuvrability, ownership, catalogue quality and encounter geometry. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate the portion addressable through operating control versus external environment. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should assign mitigation capital only where the operator can influence the outcome, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

10. Model collision probability and severity

The analysis begins with the distribution of events from harmless conjunctions to satellite loss, multi-object loss and fragmentation. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises event history, engineering vulnerability, fleet architecture, object size and relative velocity. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into expected loss plus a separately tested tail scenario. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should prevent a low average expected loss from obscuring liquidity and continuity needs.

11. Underwrite insurance coverage

Congestion becomes decision-useful when the analyst can trace launch, in-orbit, third-party liability and business-interruption protection together with exclusions and deductibles. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain policies, endorsements, limits, premiums, claims history, renewal quotes and broker market evidence. Each record needs a stable identifier and a reconciliation to the relevant

spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures recoverable cash by event and collection timing. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should recognise only enforceable cover and retain uninsured exposure in liquidity; the assumption register should show every resulting change in cash, capital and value.

12. Price self-insurance and reserves

The practical control question concerns the capital retained when commercial cover is unavailable, uneconomic or narrower than the operating risk. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes board risk appetite, fleet-loss scenarios, replacement lead times, minimum liquidity and covenant terms. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate a funded reserve or committed facility sized to restoration needs. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should approve reserve release only when exposure, insurance or replacement capacity improves, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

13. Integrate satellite reliability

The analysis begins with common-mode design faults, propulsion failures and degraded manoeuvrability that amplify congestion exposure. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises cohort telemetry, anomaly records, component lineage, production changes and failure investigations. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into joint probabilities linking reliability with collision avoidance and disposal. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should stress correlated events across one launch or hardware cohort.

14. Integrate launch and replacement lead time

Congestion becomes decision-useful when the analyst can trace the period between capacity loss, replacement manufacture, launch, orbit raising and accepted service. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain production slots, supplier lead times, launch contracts, manifests, orbital insertion and acceptance data. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures revenue interruption and bridge liquidity after a loss. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should test whether contracted replacement capacity exists when the downside occurs; the assumption register should show every resulting change in cash, capital and value.

15. Integrate end-of-life disposal

The practical control question concerns controlled deorbit, passivation, casualty risk and the probability that a failed satellite remains as a hazard. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes disposal design, residual propellant, command authority, licence conditions and actual fleet outcomes. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate cost, residual liability and renewal implications. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should fund disposal and include failures in regulatory and terminal assumptions, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

16. Price regulatory obligations

The analysis begins with licence conditions covering manoeuvrability, collision avoidance, disposal, reporting and operational coordination. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises authorisations, rulemakings, compliance reports, correspondence and waiver history. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into cash compliance cost and the value consequence of constrained operations. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should map every material obligation to an owner, date and evidence source.

17. Evaluate space-traffic coordination

Congestion becomes decision-useful when the analyst can trace the exchange of ephemerides, intent and contact data between operators and public services. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain service agreements, data-sharing protocols, response statistics and coordination test records. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures lower uncertainty, fewer unnecessary manoeuvres and faster decisions. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is

discontinuous. Decision makers should value coordination benefits only after repeatable operational evidence; the assumption register should show every resulting change in cash, capital and value.

18. Evaluate autonomous collision avoidance

The practical control question concerns software that screens, prioritises, coordinates or executes avoidance while preserving accountable control. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes validation datasets, simulation, in-orbit demonstrations, override logs, cyber controls and safety cases. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate reduced workload and response time with new model and control risks. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should stage value through verification, authority and demonstrated performance gates, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

19. Test cybersecurity and command integrity

The analysis begins with the confidentiality, integrity and availability of tracking data, manoeuvre recommendations and spacecraft commands. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises identity controls, cryptographic design, logging, incident tests, supplier access and recovery exercises. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into loss scenarios caused by corrupted data, delayed decisions or unauthorised commands. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should require independent assurance before relying on automation in value.

20. Translate congestion into usable capacity

Congestion becomes decision-useful when the analyst can trace the capacity available after manoeuvre windows, resilience reserves, disrupted beams and constrained shells. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain payload schedules, coverage maps, rerouting records, priority allocations and customer acceptance. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures billable capacity by geography, time and service class. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should value delivered capacity rather than theoretical fleet throughput; the assumption register should show every resulting change in cash, capital and value.

21. Translate congestion into customer economics

The practical control question concerns how degraded availability affects price, churn, credits, renewals and priority-service commitments. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes contracts, billing, service-level reports, customer complaints and renewal outcomes. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate contribution after congestion-related deductions. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should distinguish observed customer cash from estimated reputation or retention benefits, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

22. Construct sustainable free cash flow

The analysis begins with all recurring congestion costs, service effects, replacement capital, insurance, tax and working capital. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises reconciled financial accounts, fleet records, operating metrics and contract-level cash evidence. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into cash after maintaining safe and compliant productive capacity. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should use the same operating state in the forecast and terminal period.

23. Set a congestion-adjusted terminal value

Congestion becomes decision-useful when the analyst can trace the long-run operating burden and shell condition embedded in the period beyond the explicit forecast. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain steady-state event rates, replacement capital, regulatory outlook, capacity utilisation and return requirements. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures a terminal cash flow and multiple consistent with continuing congestion. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should reduce terminal assumptions when operating evidence cannot support stable access; the assumption register should show every resulting change in cash, capital and value.

24. Avoid double counting

The practical control question concerns overlap between operating cost, expected loss, insurance premium, replacement capital, discount rate and valuation bridge deductions. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes an assumption register showing where each risk enters the model. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate one treatment for each economic effect plus a separate tail-liquidity test. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should reconcile every adjustment before board approval, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

25. Build the hypothetical operating case

The analysis begins with a 4,200-satellite fleet with stated screening, escalation, manoeuvre and service assumptions. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises wholly hypothetical inputs identified separately from public evidence. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into annual congestion cash and life effects under a central case. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should use the case to demonstrate mechanics rather than represent an identified operator.

26. Build the hypothetical downside case

Congestion becomes decision-useful when the analyst can trace higher alert volumes, poor ephemeris quality, one significant loss, launch delay and customer credits. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain explicit management-estimated scenario variables and timing. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures cash minimum, replacement need, covenant headroom and service recovery. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should identify the earliest action date and funded response; the assumption register should show every resulting change in cash, capital and value.

27. Bridge headline to congestion-adjusted value

The practical control question concerns the deductions required when a headline case omits congestion operations, interrupted capacity and risk capital. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes transparent cash-flow and valuation adjustments with no unidentified plug. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate a board-auditable bridge from USD 8.20 billion to USD 5.05 billion. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should keep strategic technology options outside base value until milestones are achieved, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

28. Value improved tracking separately

The analysis begins with better sensors, data fusion and covariance that can reduce uncertainty and unnecessary manoeuvres. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises service-level evidence, benchmark testing and realised event outcomes. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into incremental life, workload and availability benefits after cost. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should treat the capability as an option until it is integrated and proven.

29. Value automation separately

Congestion becomes decision-useful when the analyst can trace decision support and autonomous coordination that may increase scale and response speed. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain software verification, safety assurance, cyber controls and in-orbit records. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures incremental operating savings and avoided loss after implementation and oversight. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should retain human accountability for consequential decisions; the assumption register should show every resulting change in cash, capital and value.

30. Structure financing around congestion

The practical control question concerns the match between risk maturity, replacement needs, reserves, insurance and debt service. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes lender models, covenants, collateral, committed liquidity and operating triggers. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate debt capacity based on congestion-adjusted cash flow. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should avoid financing structures that depend on unmodelled replacement or refinancing, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

31. Design transaction diligence

The analysis begins with the evidence needed to test fleet exposure, coordination capability, insurance, service consequences and liabilities. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises a data room linked by spacecraft cohort, event identifier and customer service. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into rapid reconciliation across technical, regulatory, commercial and financial workstreams. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should make high-value gaps closing conditions, price mechanisms or funded actions.

32. Establish post-close governance

Congestion becomes decision-useful when the analyst can trace the board system for orbital-risk appetite, thresholds, reserves, reporting, investment and incident review. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain approved policies, monthly dashboards, independent assurance and decision logs. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures accountable management action before operational or liquidity thresholds are breached. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should refresh the valuation when the environment or fleet evidence changes materially; the assumption register should show every resulting change in cash, capital and value.

33. Reconcile market evidence

The practical control question concerns differences in fleet maturity, shell, autonomy, launch access, insurance and customer commitments across observed companies and transactions. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes public filings, financing terms and comparable operating measures. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate a range informed by evidence rather than an unadjusted revenue multiple. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should explain adjustments before using any market observation, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

34. Incorporate tax and jurisdiction

The analysis begins with the legal entities that own spacecraft, hold licences, insure risk, bill customers and fund replacements. The point is to expose the mechanism through which orbital conditions reach customer service and cash. A fleet-wide percentage cannot show whether one shell, cohort or counterparty creates the exposure. The model should retain object, event, time and decision identifiers so the transaction team can reproduce the result.

The evidence set comprises entity-level tax bases, loss availability, transfer pricing, withholding and cash-movement rules. Records should be reconciled to the financial ledger and retained with source, date, owner and confidence. Missing data should remain a visible gap. Public catalogues can corroborate the operating environment; operator telemetry and executed contracts remain the transaction authority.

The financial model converts this evidence into cash tax and accessible liquidity by scenario. Central, downside and severe cases should change linked assumptions together, including event volume, decision time, propellant, availability, replacement and liquidity. The board should obtain transaction-specific advice before recognising tax benefits.

35. Establish evidence confidence

Congestion becomes decision-useful when the analyst can trace the quality, age, ownership and reproducibility of every material congestion assumption. The trace should distinguish observed events, engineering estimates, contractual consequences and management scenarios. It should also show which exposure the operator can control and which depends on the behaviour, tracking quality or manoeuvrability of other objects.

Diligence should obtain direct telemetry, executed contracts and regulatory records ranked above forecasts or non-binding plans. Each record needs a stable identifier and a reconciliation to the relevant spacecraft, shell, customer or cost centre. Exceptions, overrides and late changes deserve separate review because they can reveal capacity, governance or data-quality weaknesses hidden by annual averages.

Valuation then measures scenario weighting and diligence effort proportional to value exposure. The result belongs in explicit cash flow where it is recurring and in a separate tail-liquidity case where it is discontinuous. Decision makers should record what evidence would change price, financing or approval; the assumption register should show every resulting change in cash, capital and value.

36. Define the investment decision

The practical control question concerns the price, financing, conditions and operating actions supported by the congestion-adjusted evidence. It should be answered by a dated operating record rather than a narrative assertion. The assessment should show normal performance, peak workload, correlated stress and the point at which safe operation, service continuity or funding becomes constrained.

Required support includes the final value bridge, downside liquidity, consent map, reserve plan and technology milestones. The transaction team should test completeness, consistency and reproducibility, then document gaps and compensating controls. A small sample of high-consequence events should be reconstructed from first alert to final financial outcome.

The model uses the verified record to calculate a decision range with clear conditions and refresh triggers. Benefits and costs should share the same timing, tax and funding conventions. The investment committee should approve, reprice, stage or decline the transaction through an auditable record, with conditions tied to specific evidence and owners rather than a broad post-close undertaking.

Conclusion

Earth-orbit congestion is a recurring operating condition with measurable cash consequences. Conjunction screening, coordination, manoeuvres, propellant, service interruption, insurance, replacement and disposal should be connected to the same constellation cash-flow model.

The framework separates normal operations, expected loss, tail liquidity and technology options. That separation makes the value bridge auditable and prevents the same risk from appearing in operating cost, capital, insurance and the discount rate at once.

For a board or transaction team, the decisive test is whether the constellation can preserve safe access, customer service and replacement capacity through a changing orbital environment. The answer shapes price, financing, conditions and terminal value.

Appendix A. Conjunction-event register

The minimum register contains event identifier, spacecraft, shell, counterpart object, source catalogue, time of closest approach, miss distance, probability, covariance quality, alert chronology, coordination record, manoeuvre decision, delta-v, service effect, post-event assessment and accountable owner.

Appendix B. Congestion cash register

For each material event class, record analyst and system cost, propellant, payload interruption, service credit, churn exposure, replacement consequence, insurance response, regulatory action and tax treatment. This register connects orbital evidence with cash.


Appendix C. Investment checklist

The approval file should contain fleet and shell exposure, tracking and coordination quality, manoeuvre authority, propellant margins, service impact, insurance, replacement lead time, disposal, liquidity, downside actions and technology milestones. Evidence should be dated and owned.

Appendix D. Worked-case figures and tables


Figure 1. Congestion-adjusted constellation valuation architecture

Orbital congestion to enterprise value




Proposed architecture connecting orbital environment, conjunction operations, customer capacity, cash and value.

Figure 2. Hypothetical annual conjunction funnel




Hypothetical screened messages, escalated events, coordinated events and avoidance manoeuvres.

Figure 3. Hypothetical annual congestion cash bridge




Hypothetical customer contribution less recurring orbital-risk costs and sustaining capital.

Figure 4. Hypothetical enterprise-value bridge



Hypothetical bridge from headline enterprise value to congestion-adjusted value and a net technology option.

Figure 5. Hypothetical orbital-congestion risk heat map



Illustrative probability and cash-consequence assessment for key operating risks.

Table 1. Orbital-congestion evidence register

Evidence	Required measure	Valuation use
Fleet and shell	active degraded and non-manoeuverable objects	exposure perimeter
Conjunction ledger	alerts probability covariance and decisions	workload and event risk
Manoeuvre record	delta-v time and service effect	life and capacity
Insurance	cover exclusions limits and collection	recoverable cash
Replacement	production launch and acceptance lead time	interruption and liquidity

Proposed diligence fields.

Table 2. Hypothetical annual event funnel

Stage	Annual volume	Primary economic effect
Screened messages	68,000	systems and analyst workload
Escalated events	2,400	engineering review
Coordinated events	920	operator communication and scheduling
Avoidance manoeuvres	410	propellant capacity and life
Material service events	24	credits contribution and churn

All figures are illustrative and describe no identified operator.

Table 3. Congestion cash controls

Effect	Evidence	Cash treatment
Screening and operations	payroll vendor and computing records	recurring operating cost
Manoeuvre and recovery	command telemetry and service records	operating and life effect
Satellite loss	incident and replacement evidence	expected loss and capital
Insurance recovery	policy and accepted claim	timed recoverable cash
Tail event	correlated scenario and restoration plan	liquidity test

Proposed reconciliation.

Table 4. Hypothetical valuation bridge

Step	Amount	Evidence required
Headline enterprise value	8.20	forecast and market assumptions
Recurring congestion operations	minus 0.55	sustainable operating cost
Capacity interruption	minus 0.45	contract and service evidence
Shortened life and replacement	minus 0.70	propulsion and cohort model
Tail risk and liquidity	minus 0.65	severe case and funded response
Terminal adjustment	minus 0.80	steady-state shell assumptions
Congestion-adjusted value	5.05	integrated model
Automation and tracking option	plus 0.55	performance milestones
Implementation cost	minus 0.20	executable programme budget
Final illustrative value	5.40	board-approved evidence set

All amounts are illustrative USD billions.

Table 5. Orbital-risk approval gate

Gate	Approval evidence	Stop condition
Environment	reconciled shell and object exposure	material untracked exposure
Operations	tested screening and response	unsafe workload or latency
Service	measured interruption and recovery	unpriced material credits
Capital	funded replacement and reserve	liquidity gap before restoration
Regulation	licences reports and disposal plan	unresolved operating restriction
Governance	authority cyber and assurance	unclear consequential accountability

Proposed board control.

Table 6. Financing structure by uncertainty

Requirement	Potential instrument	Evidence gate
Automation development	equity	validated safety case
Tracking integration	vendor or partner finance	accepted service levels
Routine fleet operations	corporate facility	stable congestion-adjusted cash
Contracted capacity	asset or project-style debt	service and replacement reserve
Tail restoration	committed liquidity	severe-case restoration plan

Table 7. Board reporting dashboard

Dimension	Core measure	Trigger
Environment	conjunctions by shell and object class	event-rate or data-quality deterioration
Decisions	alerts escalation and response time	authority or workload breach
Fleet	delta-v life degradation and disposal	replacement acceleration
Customers	availability credits churn and contribution	negative service economics
Insurance	limits exclusions claims and renewal	unfunded exposure
Liquidity	reserve replacement and covenant headroom	restoration shortfall
Governance	overrides incidents and assurance	control failure

Proposed monthly decision record.

Frequently asked questions

Q: Why should orbital congestion enter valuation directly? A: It creates recurring operating cost, consumes propulsion, can interrupt service, accelerates replacement and creates uninsured tail exposure. Each effect can change cash flow or terminal value.

Q: Is collision probability alone sufficient? A: No. The model also needs alert volume, data quality, manoeuvre rules, service effects, replacement lead time, insurance and liquidity.

Q: How should avoidance manoeuvres be valued? A: Use observed operating cost, delta-v, life effect, payload interruption and customer consequences by spacecraft cohort and shell.

Q: How should insurance enter the model? A: Recognise enforceable cover after exclusions, deductibles, limits and collection timing. Retain uninsured loss and restoration funding in the downside case.

Q: Does automation automatically create value? A: Value requires verified performance, safety assurance, cybersecurity, regulatory acceptance and realised operating or loss reduction after implementation cost.

Q: How should tail events be treated? A: Keep recurring expected loss in cash flow and test severe correlated events separately for liquidity, service continuity and covenant headroom.

Q: What belongs in terminal value? A: The terminal period should include steady-state congestion operations, service effects, replacement capital, regulatory compliance and a shell condition consistent with available evidence.

Q: What is the central board decision? A: The board must decide whether congestion-adjusted cash, restoration liquidity, operating controls and verified options support the proposed price and financing.

References

- [1] European Space Agency, ESA Space Environment Report 2025.
https://www.esa.int/Space_Safety/Space_Debris/ESA_Space_Environment_Report_2025
- [2] European Space Agency, ESA Space Environment Report 2026.
https://www.sdo.esoc.esa.int/publications/Space_Environment_Report_I10R0_20260501.pdf
- [3] NASA, Conjunction Assessment Risk Analysis. <https://www.nasa.gov/conjunction-assessment/>
- [4] NASA, NPR 8079.1, NASA Spacecraft Conjunction Analysis and Collision Avoidance for Space Environment Protection.
<https://nodis3.gsfc.nasa.gov/displayDir.cfm?t=NPR&c=8079&s=1>
- [5] Federal Communications Commission, Mitigation of Orbital Debris in the New Space Age, Second Report and Order, FCC 22-74.
<https://docs.fcc.gov/public/attachments/FCC-22-74A1.pdf>

- [6] Federal Communications Commission, SpaceX Gen2 authorisation and orbital-debris conditions, FCC 22-91. <https://docs.fcc.gov/public/attachments/FCC-22-91A1.pdf>
- [7] European Space Agency, CREAM: avoiding collisions in space through automation. https://www.esa.int/Space_Safety/Space_Debris/CREAM_avoiding_collisions_in_space_through_automation
- [8] European Space Agency, Space Debris Mitigation Requirements, ESSB-ST-U-007. <https://technology.esa.int/upload/media/ESA-Space-Debris-Mitigation-Requirements.pdf>
- [9] United Nations Office for Outer Space Affairs, Guidelines for the Long-term Sustainability of Outer Space Activities. <https://www.unoosa.org/oosa/en/ourwork/topics/long-term-sustainability-of-outer-space-activities.html>
- [10] Inter-Agency Space Debris Coordination Committee, Space Debris Mitigation Guidelines. https://www.iadc-home.org/documents_public/view/id/82
- [11] International Organization for Standardization, ISO 24113 Space systems: Space debris mitigation requirements. <https://www.iso.org/standard/83474.html>
- [12] U.S. Space Command, Space-Track conjunction data and operator services. <https://www.space-track.org/documentation>
- [13] NASA, Orbital Debris Program Office. <https://orbitaldebris.jsc.nasa.gov/>
- [14] NASA, Orbital Debris Mitigation Standard Practices. https://orbitaldebris.jsc.nasa.gov/library/usg_orbital_debris_mitigation_standard_practices_november_2019.pdf
- [15] Astroscale Holdings, Annual Report 2025. <https://ir.astroscale.com/en/library/annual.html>
- [16] Satellogic, 2024 Annual Report. <https://www.sec.gov/Archives/edgar/data/1874315/000162828025014951/satl-20241231.htm>
- [17] Spire Global, 2025 Annual Report. <https://www.sec.gov/Archives/edgar/data/1816017/000119312526116169/spir-20251231.htm>
- [18] Globalstar, 2024 Annual Report. <https://www.sec.gov/Archives/edgar/data/1366868/000136686825000033/gsat-20241231.htm>
- [19] IFRS Foundation, IAS 16 Property Plant and Equipment. <https://www.ifrs.org/issued-standards/list-of-standards/ias-16-property-plant-and-equipment/>
- [20] IFRS Foundation, IAS 36 Impairment of Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>
- [21] IFRS Foundation, IFRS 13 Fair Value Measurement. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>
- [22] International Private Equity and Venture Capital Valuation Guidelines. <https://www.privateequityvaluation.com/>
- [23] International Valuation Standards Council, International Valuation Standards. <https://www.ivsc.org/standards/>
- [24] European Space Agency, Space environment health index 2025. https://www.esa.int/ESA_Multimedia/Images/2025/10/Space_environment_health_index_2025
- [25] International Telecommunication Union, Non-geostationary satellite orbit licensing framework toolkit. <https://www.itu.int/hub/publication/d-pref-ngso-2026/>

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and closure.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners, he is Managing Partner and leads the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.