

**FLEXIBLE SPACE JUNK ALLOCATION & WASTE ABATEMENT
(SPACE JAWA): FLEXIBILITY FRAMEWORK TO SCREEN
STRATEGIES & OPTIONS FOR SUSTAINABLE ON-ORBIT
SERVICING INFRASTRUCTURES IN LEO**

A Dissertation
Presented to
The Academic Faculty

By

Madeline Bowne

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Engineering
Department of Aerospace Engineering

Georgia Institute of Technology

December 2025

© Madeline Bowne 2025

**FLEXIBLE SPACE JUNK ALLOCATION & WASTE ABATEMENT
(SPACE JAWA): FLEXIBILITY FRAMEWORK TO SCREEN
STRATEGIES & OPTIONS FOR SUSTAINABLE ON-ORBIT
SERVICING INFRASTRUCTURES IN LEO**

Thesis committee:

Dr. Dimitri Mavris
Aerospace Engineering
Georgia Institute of Technology

Dr. Koki Ho
Aerospace Engineering
Georgia Institute of Technology

Dr. Mariel Borowitz
Sam Nunn School of International Affairs
Georgia Institute of Technology

Dr. Mark Whorton
Aerospace Engineering
NASA Marshall Space Flight Center & Georgia Institute of Technology

Dr. Olivia Fischer
Aerospace Engineering
Georgia Institute of Technology

Dr. Tristan Sarton du Jonchay
Aerospace Engineering
Georgia Institute of Technology

Date approved: November 7th, 2025

“There is no real direction here, neither lines of power nor cooperation. Decisions are never really made – at best they manage to emerge, from a chaos of peeves, whims, hallucinations and all around a**holery.”

Thomas Pynchon, Gravity’s Rainbow

ACKNOWLEDGMENTS

There are many people I'd like to thank for their endless support. A thesis would be a long and lonely endeavor were it not for the people who support you along the way.

I'd like to thank my advisor, Dr. Mavris, for his mentorship and support. One of the greatest lessons I've taken from the thesis process is the art of clear argumentation and precision in communication. In this day and age, it has never been more important to leverage verifiable information and logic to move a conversation forward, and to do so in a manner that is accessible, captivating, and digestible to a general audience. Dr. Mavris has taught his students not only how to identify research opportunities and propose new methods to tackle problems, but how to do so through accessible storytelling. I still have a lot to learn and will continue striving to improve. Dr. Mavris instills a lifelong dedication to problem-solving, substantiation, and effective technical communication in all of his students.

I would also like to thank my technical advisor, Dr. Tristan Sarton du Jonchay, who has expressed tremendous support and excitement in my thesis topic. He has been instrumental in bringing my CAAS concept to a larger scale—launching a grand challenge related to CAAS and inspiring another PhD thesis related to the topic. His work spanning terrestrial and space logistics has been a fantastic example of how to draw inspiration across fields. His advice and feedback have been instrumental in helping me hone the CAAS CONOPs and focus the scope of this dissertation.

I'd also like to extend a thank you to my other committee members as well. Dr. Olivia Fischer, Dr. Koki Ho, Dr. Mariel Borowitz, and Dr. Mark Whorton have all provided helpful feedback. This thesis is made possible thanks to their support.

I got by with a little help from my friends. From late nights in the “cube farm” our first year to months-long study sessions for the PhD qualifying exams to regular

thesis meetings that kept us accountable, my friends at ASDL have provided academic and emotional support that made my grad school career not only manageable, but thoroughly enjoyable. Over the last five years, I've had the privilege of watching you all grow as researchers and as people eager to take on the world's problems while supporting your communities. Outside the lab, we've had so much fun exploring Atlanta together, from karaoke at Dark Horse Tavern to long strolls down the Beltline. I'm sad this chapter is closing, but I'm excited to see where life takes you all. Thank you to my ASDL friends: Leon Chen, Sissy Pastra, Xi Wang, Eleni Sotiropoulos-Georgopoulos, and Ezgi Balkas.

And to my friends back home in NJ – the Little Falls Fam – thank you for all your support, encouragement, and friendship during my time at GT. In the last few years, between grad school and internships, I've lived in all four corners of the United States, but NJ – and the Little Falls Fam – has always been home. You have been a balancing influence in my life and I appreciate all the love and support you provide.

To my parents, Mary Jane and Walter, and my sister, Nancy: thank you for your inspiration and support throughout my entire life. You have always believed in me and my capability to tackle any challenge head-on. Each of you has an unrelenting commitment to your respective crafts. You approach your interests with such vim and vigor, demonstrating intellectual curiosity and a hunger for new understanding. In so many ways, you model what it means to be lifelong students, which is something I hold near and dear to my heart. Thank you all so much for your guidance through this process. Thank you for having so much fun with my JAWA acronym and gifting me all the Jawa-themed desk toys. But most importantly, thank you for always standing by me, through both my successes and my failures.

Last but not least, I'd like to thank my fiancé, Brian Lange. He's a cool dude.

...

Didn't think you'd get away that easily, did you? Thank you, Brian, for being the circulation to my airfoil and the gravity-gradient stabilization to my spacecraft. I realize that if you were to write these acknowledgments, there would be way more puns, but I'll do my best. You have been unceasingly supportive, helpful, and patient. You've been along for this whole emotional rollercoaster that is the PhD experience, celebrating my wins and encouraging me to get back up after my losses. You seem to know, better than I do, why I decided to pursue a PhD, and you have never let me lose sight of myself. At the same time, you provide a grounding influence, helping me find joy in this period of my life. The thesis experience is unique for a young person getting started in their career; it brings unprecedented freedom and flexibility to dive deep into a topic of their choice. At the same time, this can be a very painful experience, as we wade around in a soup of abstraction with no clear answers and no clear directions. You have been a compass for me through all of this, helping me remember why I even started. Aside from your moral support, you have been tremendously helpful in the construction of this thesis itself—giving me feedback on the flow of my presentations, proofreading my thesis document and leaving helpful (and often humorous) comments. You, and your entire family, have been remarkably encouraging through this process. Thank you, Brian Lange, for your unwavering love and support and being my L(agr)ange point.

TABLE OF CONTENTS

Acknowledgments	iv
List of Tables	xiii
List of Figures	xvi
List of Acronyms	xx
Summary	xxvi
Chapter 1: Introduction	1
1.1 Thesis Objective & Contributions	2
1.2 Thesis Organization	4
Chapter 2: Background and Motivation	8
2.1 Growth of the Commercial Space Industry	8
2.2 Orbital Congestion & Space Debris	9
2.3 Global Response to the Space Debris Problem	10
2.4 Environmental Impact	15
2.5 Circular Space Economies	20
2.6 Current State of OSAM/OOS	23
2.6.1 OOS Business Case	25

2.7	Developments in LEO-based OOS	28
2.8	Hypothesis 1: Collection Hubs	28
2.8.1	LEO-Specific OOS Challenges	28
2.8.2	Improving OOS CONOPs for LEO	29
2.8.3	Understanding the Needs of Mega-Constellations	38
2.8.4	Graveyard Orbits and the Value of Temporary Abandonment .	42
2.8.5	Enabling Technology for Earth Return	44
2.8.6	Gaps and Proposed Solutions	49
2.9	Hypothesis 2: Flexibility Frameworks	53
2.9.1	Flexibility Frameworks	53
2.9.2	Real World Examples	55
2.9.3	Flexibility Frameworks for OOS	59
2.9.4	Flexible OOS Concepts	66
2.10	Hypothesis 3: Policy	71
2.10.1	The Policy Problem: Environmental Impact and Market Failures	71
2.10.2	Current Policy Landscape and Regulatory Gaps	72
2.10.3	Policy Mechanisms to Mitigate Orbital Congestion	75
2.10.4	Policy Mechanisms from Analogous Industries	82
2.10.5	Policy Design for Closing the Business Case Gap	90
2.10.6	Research Gaps and Framework Development	94
2.11	Thought-Experiment and Real-World Analogies	97
2.12	Overarching Research Framework	99

Chapter 3: Methodology	102
3.1 CONOPs Summary	102
3.1.1 Core Infrastructure Components	105
3.1.2 Operational Framework	107
3.1.3 Collection and Processing Options	109
3.1.4 Flexibility and Adaptability	111
3.1.5 Novel CONOPs	112
3.2 Problem Properties and Framework Requirements	114
3.3 Formulation Question 1: Framework Selection	115
3.3.1 Formulation Question 1.1: Baseline Establishment	117
3.3.2 Formulation Question 1.2: Uncertainty Variable Selection and Characterization	122
3.3.3 Formulation Question 1.3: Concept Generation	136
3.3.4 Formulation Question 1.4: Design Space Exploration	138
3.4 Formulation Question 2: Simulation Method Selection and Development	143
3.4.1 Space Logistics Modeling	144
3.4.2 Discrete Event Simulation	146
3.5 Formulation Question 3: Decision Rule Creation and Calibration . . .	148
3.5.1 Experimental Tuning Process	161
3.6 Formulation Question 4: DES Component Modeling	162
3.6.1 Formulation Question 4.1: Object Modeling	162
3.6.2 Formulation Question 4.2: Cost Modeling	176
3.6.3 Formulation Question 4.3: Trajectory Modeling	182

3.6.4	Formulation Question 4.4: Emissions Modeling	184
3.6.5	Formulation Question 4.5: Policy Modeling	186
3.6.6	Addressing Regulatory Gaps	192
3.6.7	Policy Impact Integration in Cost-Benefit Analysis	193
3.6.8	Decision Logic Integration	195
3.6.9	Fund Flow Dynamics	195
3.7	Framework Integration and Experimentation	196
Chapter 4: Screening Flexible Options & Experimentation		200
4.1	Experimentation Overview	200
4.1.1	Use Case	203
4.1.2	Testing for Statistical Significance	203
4.1.3	Sensitivity Testing	209
4.2	H1 Experimentation: Collection as a Service	211
4.2.1	Configuration Set	215
4.2.2	VARG Results for Experiment 1	215
4.2.3	Experimental Support for Hypothesis 1	226
4.2.4	Hypothesis 1 Substantiation	231
4.3	H2 Experimentation: Flexibility	233
4.3.1	Configuration Set	235
4.3.2	VARG Results for Experiment 2	237
4.3.3	Sensitivity Testing	245
4.3.4	Experimental Support for Hypothesis 2	254

4.4	H3 Experimentation: Policy	258
4.4.1	Experiment 3.1: Policy Parameter Down-Selection Analysis . .	260
4.4.2	Experiment 3.2: Comparison and Analysis for Down-Selected Policies	266
4.4.3	Experiment 3.3: Policy and Flexibility for Configurations with Greater Than 1 Initial Depot	275
4.4.4	Experimental Support for Hypothesis 3	283
4.4.5	Hypothesis 3 Substantiation	286
4.5	Experimental Support for Overarching Hypothesis	288
4.5.1	Major Findings from Experiment 1	289
4.5.2	Major Findings from Experiment 2	292
4.5.3	Major Findings from Experiment 3	292
4.5.4	Minor Supporting Findings	293
4.5.5	Integration: Addressing the Overarching Research Question .	295
4.5.6	Implications for Circular Space Economy Development and Closing Remarks	296
Chapter 5: Conclusion	299
5.1	Major Contributions and Key Findings	299
5.1.1	Non-Intuitive Findings	301
5.2	Critical Insights for Decision-Makers	303
5.3	Actionable Recommendations	305
5.3.1	For Policy Makers	305
5.3.2	For Satellite Constellation Operators	308
5.3.3	For Satellite Manufacturers	310

5.3.4 For On-Orbit Servicing Providers	311
5.4 Lessons Learned and Recommendations for Future Researchers	313
5.5 Future Work	314
5.6 Closing Perspective	318
Appendices	320
Appendix A: Assumptions, Equations, and Models	321
Appendix B: Decision Rule Algorithms	336
References	341
Vita	359

LIST OF TABLES

2.1	Highlighted Constellation Projections	9
2.2	OOS/ADR Companies: Past, Present, and Future Missions with RPOD/RPO Capabilities	24
2.3	Tradespace Scenarios [3]	31
2.4	Spacecraft Retrieved by Space Shuttle and Returned to Earth	46
2.5	Space Mission Options [10]	67
2.6	Morphological Matrix for OOS/ISAM Part 1: Architecture and Capabilities	69
2.7	Morphological Matrix for OOS/ISAM Part 2: Operations and Flexible Options	70
3.1	Use Case Simulation Parameters	118
3.2	Sources of Uncertainty and Modeling Parameters	119
3.3	Latin Hypercube DOE Parameters for Uncertain Variables	122
3.4	Summary of Simulation Costs	176
3.5	Initial Cost Assumptions (Year: 2025)	178
4.1	Space Infrastructure Configuration Parameters and Dependencies . .	214
4.2	Configuration Combination Summary	214
4.3	Parameter Settings for Top Performing Configurations: Experiment 1	217

4.4	Experiment 1: Total Cost	221
4.5	Experiment 1: Total NOx Emissions	222
4.6	Experiment 1: Number of Refurbishments	223
4.7	Key Changes from Previous Configuration Logic	235
4.8	Configuration Combination Summary	236
4.9	Configuration Parameter Settings for All Depots - Flexible	238
4.10	Experiment 2: Total Cost	240
4.11	Experiment 2: Total NOx Emissions	241
4.12	Experiment 2: Number of Refurbishments	242
4.13	Sensitivity Test Configuration Parameters	246
4.14	Sensitivity analysis results for key uncertain variables and aspects . .	247
4.15	Deterministic sensitivity analysis results for operational cost parameters (Part 1)	252
4.16	Deterministic sensitivity analysis results for operational cost parameters (Part 2)	253
4.17	Policy Configuration Parameters	261
4.18	Experiment 3.1 Down-Selection Summary	265
4.19	Infrastructure Configurations for Policy Comparison Experiment 3.2 and 3.3	266
4.20	Top Performing Configuration Parameter Settings - Experiment 3.2 .	268
4.21	Experiment 3.2: Total Cost	272
4.22	Experiment 3.2: Total NOx Emissions	273
4.23	Experiment 3.2: Number of Refurbishments	274
4.24	Downselected Parameter Configurations for Policy Comparison Experiment 3.3	275

4.25	Top Performing Configuration Parameter Settings - Experiment 3.3	276
4.26	Experiment 3.3: Total Cost	280
4.27	Experiment 3.3: Total NOx Emissions	281
4.28	Experiment 3.3: Number of Refurbishments	282
A.1	Use Case Simulation Parameters	321
A.2	Rocket Performance Parameters	322
A.3	Sources of Uncertainty and Modeling Parameters	322
A.4	Initial Cost Assumptions (Year: 2025)	325
A.5	Latin Hypercube DOE Parameters for Uncertain Variables	329

LIST OF FIGURES

1.1	Space Junk Allocation & Waste Abatement (JAWA)	5
1.2	Thesis Logic Diagram	6
1.3	Thesis Methodology	7
2.1	Cumulative Objects in Space [23]	10
2.2	Visualization of Tracked Objects in LEO. Active satellites are blue, large debris is red, small debris is gray [24]	11
2.3	3 Phases of the Yap et al. Socio-Technical Configuration Analysis show Sustainability Logic (purple) shift to a central focus by 2016 (Phase 3) [26]	13
2.4	Elements of a Circular Space Economy	21
2.5	Circular Economies in Space [25]	22
2.6	OOS Pod CONOPs [3]	30
2.7	Orbital Phasing [3]	32
2.8	Depot Drift Time vs. Orbit and Delta-V [95]	32
2.9	Cooperative Manuevering in LEO [96]	34
2.10	Multi-Level Spare Strategy	40
2.11	Inventory Spare Strategy [14]	41
2.12	Illustration of the CAAS Warehouse Concept	52
2.13	Flexibility Frameworks at Work	57

2.14	Value of Compound Abandonment [142]	64
2.15	GAO Policy Framework [150]	76
2.16	Real World Analogies of Circular Systems	98
2.17	Framework Gap for Experimentation Conditions	101
3.1	Comparison of sparing strategies for satellite resupply	106
3.2	CAAS Upgrade Option	107
3.3	ADR Launched from Earth	108
3.4	ADR Deployed from Parking Orbit	109
3.5	Plane Resupply with Rendezvous Option	110
3.6	Refurbishment on Earth Option	111
3.7	Flexibility Framework Phases [6]	116
3.8	Launch Cost Predictions [198]	124
3.9	Cone of Uncertainty, Geometric Diffusion [142]	125
3.10	Spacecraft Present Value with Technology Obsolescence [201]	129
3.11	Monte Carlo Scenario Construction	136
3.12	Example of a VARG Plot (Profit)	143
3.13	Satellite-Level Decision Tree Logic	150
3.14	RPO Upgrade Decision Rule Tree	154
3.15	Satellite Refuelability/Repairability Decision Tree	155
3.16	Add Warehouse Decision Rule Tree	157
3.17	Upgrade Warehouse Decision Rule Tree	159
3.18	Top-Level Methodology Diagram	163

3.19 Methodology Logical Flow Diagram	175
3.20 Use of the JAWA Framework	197
3.21 Experimentation Logic Flow Diagram	198
3.22 Experimental Method Diagram	199
4.1 The Bavaria	200
4.2 Supporting Conditions for Hypotheses	202
4.3 Experiment 1 Methodology to Support Hypothesis 1	212
4.4 Total Cost for Experiment 1: VARG Plot	218
4.5 Total NOx Emissions for Experiment 1: VARG Plot	219
4.6 NOx Emissions vs. Year of Earth-Return Maturity	230
4.7 Experiment 2 Methodology to Support Hypothesis 2	234
4.8 Total Cost for Flexible CAAS vs. Inflexible CAAS vs. Baseline for Experiment 2	239
4.9 Experiment 1 Methodology to Support Hypothesis 3	259
4.10 Experiment 3.2: Total Cost VARG Plot	270
4.11 Experiment 3.2: Total NOx Emissions VARG Plot	270
4.12 Pareto Frontier: Total NOx Emissions (kg) vs. Total Costs (\$)	278
4.13 Pareto Frontier: Total Satellite Refurbishments vs. Total Costs (\$)	279
B.1 Appendix: Satellite-Level Decision Tree Logic	336
B.2 Appendix: Satellite Refuelability/Repairability Decision Tree	337
B.3 Appendix: RPO Upgrade Decision Rule Tree	338
B.4 Appendix: Add Warehouse Decision Rule Tree	339

LIST OF ACRONYMS

ABMS	Agent-Based Modeling and Simulation
ACP	Aggregate Collision Probability
ADR	Active Debris Removal
a	Semi-major axis of the orbit
a_{transfer}	Semi-major axis of the transfer orbit
α_m	Drift [1/time]
ANOVA	Analysis of Variance
ASAT	Anti-Satellite Weapon
β	Collision hazard rate per satellite per year [1/year]
β	Weibull distribution parameter [dimensionless]
BC	Black Carbon
CAAS	Collection-as-a-Service
CAFE	Corporate Average Fuel Economy
CAPEX	Capital Expenditure
C	CAAS configuration shorthand
C_{base}	Base collision cost [\$10,000]
$C_{\text{collision}}$	Collision cost [\\$]
C_{failure}	Cost of warehouse failure [\\$]
CO₂	Carbon Dioxide
CONOPs	Concept of Operations
COPUOS	UN Committee on the Peaceful Uses of Outer Space
CORE	Common Orbital Refueling Elements
COSMIC	Cleaning Outer Space Mission through Innovative Capture
CP	Chemical Propulsion

<i>D</i>	Depot/Warehouse configuration shorthand
DARPA	Defense Advanced Research Projects Agency
DES	Discrete Event Simulation
DOE	Design of Experiments
DOF	Degrees of Freedom
Δv	Change in velocity (general) [km/s]
$\Delta v_{\text{Hohmann}}$	Total delta-V required for a Hohmann transfer [km/s]
Δv_{total}	Total delta-V for a transfer [km/s]
$\Delta\Omega$	Change in RAAN [angle]
\dot{m}	Mass flow rate [kg/s]
$\dot{\Omega}$	Rate of change of RAAN due to J2 perturbation [deg/s or rad/s]
<i>e</i>	Eccentricity [dimensionless]
ECLSS	Environmental Control and Life Support System
ELSA	End-of-Life Services by Astroscale
EOL	End-of-Life
EP	Electric Propulsion
EPA	Environmental Protection Agency
ESA	European Space Agency
EU ETS	European Union Emissions Trading System
<i>f</i>	Current failure rate [failures/year]
f_0	Original/baseline failure rate [failures/year]
<i>F</i>	Thrust force [N]
FAA	Federal Aviation Administration
FCC	Federal Communications Commission
FINE	Fine policy configuration
F-statistic	Ratio of explained to unexplained variance
<i>g</i>	Gravitational acceleration (standard) [m/s ²]

g_0	Standard gravitational acceleration at Earth's surface [9.80665 m/s ²]
GEO	Geostationary Earth Orbit
GRIP	Grappling and Resupply Interface for Products
HCSC	Health Case Services Corporation
HTF	Highway Trust Fund
i	Orbital inclination [deg or rad]
IADC	Inter-Agency Space Debris Coordination Committee
INIT_SUB	Policy scheme initial subsidy
ISAM	In-Space Servicing Assembly and Manufacturing
I_{sp}	Specific impulse [s]
ISS	International Space Station
ISO	International Organization for Standardization
J_2	Earth's second zonal harmonic coefficient [1.08263×10 ³ , dimensionless]
JAWA	Junk Allocation and Waste Abatement
k	Weibull shape parameter
λ	Learning exponent or rate parameter
λ	Weibull scale parameter [years]
LEO	Low Earth Orbit
LHS	Latin Hypercube Sampling
m_0	Initial mass [kg]
m_{prop}	Propellant mass [kg]
mass_{dry}	Mass of the vehicle without propellant (dry mass) [kg]
mass_{wet}	Total mass of the vehicle including propellant (wet mass) [kg]
MAUT	Multi-Attribute Utility Theory
MEO	Medium Earth Orbit
MEV	Mission Extension Vehicle
μ	Mean of log-normal distribution [dimensionless]

μ	Standard gravitational parameter (Earth) [3.986*10e5 km^3/s^2]
μ_{launch}	Mean launch delay [months]
MILP	Mixed Integer Linear Programming
MINLP	Mixed Integer Nonlinear Programming
MRV	Mission Robotic Vehicle
MSS	Mobile Satellite Services
MTBF	Mean Time Between Failure
n	Mean motion of the orbit [rad/s or deg/day]
N_{fail}	Number of failed satellites
N_{max}	Maximum number of units considered
NEPA	National Environmental Policy Act
NO_x	Nitrogen Oxides
NPV	Net Present Value
OCD	Orbital Cleaning Drones
OOS	On-Orbit Servicing
OPEX	Operational Expenditure/Costs
OSAM	On-Orbit Servicing and Manufacturing
OUF	Orbital Use Fee
P	Policy configuration shorthand
P_{cascade}	Probability of cascade collision event
$P_{\text{collision}}$	Probability of collision
P_{min}	Minimum collection success probability level (asymptotic limit)
$p_{\tau}^{(m)}(x)$	Log-normal probability density function
POD-U	Proximity Operations and Docking-Undocking
POLICY_EN	Policy scheme configuration (enabled/disabled)
PPP	Polluter-Pays Principle
PREM	Annual premium policy configuration

p-value	Statistical significance indicator
r	Risk-free interest rate [%]; Annual reduction rate [%/year]
r_0	Radius of the initial orbit [km or m]
r_1	Radius of the target orbit [km or m]
RAAN	Right Ascension of Ascending Node
RAFTI	Rapidly Attachable Fluid Transfer Interface
R_E	Earth's equatorial radius [6378.137 km]
REBATE	Rebate for refurbishment policy configuration
REFUND	Refund condition policy configuration
ReHAB	Recycling Hub and Base
Rf	Refuelable configuration shorthand
RGGI	Regional Greenhouse Gas Initiative
ROA	Real Options Analysis
RPO	Rendezvous and Proximity Operations
RPOD	Rendezvous, Proximity Operations, and Docking
RSGS	Robotic Servicing of Geosynchronous Satellites
S_0	Initial collection probability
SBIR	Small Business Innovation Research
σ	Standard deviation of log-normal distribution
σ_m	Volatility parameter [$1/\sqrt{time}$]
SME	Subject Matter Expert
SSA	Space Situational Awareness
STCA	Socio-Technical Configuration Analysis
t	Time in years [years]
t_{burn}	Time required for low thrust maneuver
t_{failure}	Time to failure
t_{RAAN}	Time required for RAAN drift due to J2 precession

t_{transfer}	Total transfer time
τ	Time parameter in stochastic models [years]
TAX_PCT	Tax percentage policy configuration
TAX_SHAPE	Tax shape parameter policy configuration
T_{cascade}	Time parameter for cascade events [years]
θ	Plane change angle
$\theta_{\text{obs},i}$	Obsolescence intensity parameter
$T_{\text{obs},i}$	Time to technology obsolescence [years]
$T_{\text{processing}}$	Processing time
TRL	Technology Readiness Level
$U(a, b)$	Uniform distribution between a and b
$u_i(t)$	Utility function at time t
UKSA	UK Space Agency
UMPIRE	Universal Mission Planner to Investigate Refueling Effectiveness
UN	United Nations
UNOOSA	UN Office for Outer Space Affairs
Upg	Upgraded configuration shorthand
$u_{\text{o},i}$	Initial utility value
u_{total}	Total utility
v_0	Circular orbital velocity at the initial altitude [km/s]
v_1	Circular orbital velocity at the target altitude [km/s]
$v_{\text{transfer},0}$	Velocity at the start of the transfer orbit [m/s or km/s]
$v_{\text{transfer},1}$	Velocity at the end of the transfer orbit [km/s]
VaG	Value at Gain (95th percentile)
VaR	Value at Risk (5th percentile)
VARG	Value at Risk/Value at Gain

SUMMARY

Between the start of my PhD thesis process in January 2023 and its defense, 1420 objects, including rocket bodies, payloads, and tracked space debris, have re-entered Earth's atmosphere [1]. Over the next decade, as the commercial space industry continues to grow and mega-constellations fully deploy, this rate of reentry is likely to increase tremendously. Launch costs are low and getting lower, contributing to the popularity of proliferated constellations of cheap, short-lived satellites in Low Earth Orbit (LEO). These constellations provide invaluable services, such as Earth imagery for climate science and internet for under-served regions of the globe. However, growing orbital traffic, particularly in LEO, contributes to greater collision risk. To mitigate orbital congestion, deorbiting LEO satellites is now mandated by FCC licensure [2].

There's a good chance that Earth's built-in garbage disposal system isn't free, and in several years, we could be paying the environmental cost. While scientists recognize that re-entering space debris could damage the ozone and impact climate change, the extent of which is unknown. It's wise for mega-constellation operators to be mindful of the environmental uncertainty since future policies or social pressure could influence their operations.

Transitioning from a linear economy of single-use satellites to a circular economy of re-use could mitigate this source of uncertainty, as well as improve profits for spacecraft operators in the long run. A critical enabler for circular space economies is On-Orbit Servicing (OOS), a concept that is far more popular for Geosynchronous orbit (GEO) than it is for LEO. OOS in LEO is less appealing due to decreasing launch and satellite cost and inefficient out-of-plane maneuvering. There have been some concept proposals to improve OOS for LEO, such as proliferated servicing pods [3], cooperative maneuvering [4], and focusing on scheduled service instead of on-

demand. However, other interventions and strategies will be necessary to sufficiently motivate the private sector to invest in OOS for LEO.

Policy and regulations could incentive OOS in LEO by either penalizing or rewarding spacecraft operators for their disposal and servicing choices. Market-based policies internalize the cost of environmental damage, influencing private companies to opt for environmentally-conscious decisions [5]. While there are currently no policy proposals to encourage LEO-based OOS, there are market-based policy proposals regarding orbital congestion. Given the close relationship between orbital debris mitigation/remediation and On-Orbit servicing, these orbital congestion policies could simultaneously drive demand for OOS with little to no adjustment.

On-orbit servicing is capital-intensive, operates over a long timeline, and is fraught with various sources of uncertainty; qualities that many infrastructure projects share. Flexibility frameworks have successfully identified value in previous infrastructure efforts, such as nuclear power plants, offshore oil rigs, waste-to-energy facilities, water reservoirs, and many others. The study of flexibility is interested in improving engineering system value via novel, flexible design methods and procedures, referred to as enablers and strategies, allowing decision-makers to improve economic value, sustainability, and resilience amidst an uncertain landscape [6]. In their publications and theses, Saleh, Lamassoure, and Hastings [7] [8] [9] [10] provide insight on the flexibility value of On-Orbit Servicing. On-orbit servicing provides more than the possibility for cost savings, it provides a source of flexibility for spacecraft operators [9]. Their framework focuses on the customer point of view, which is valuable for understanding the demand for On-Orbit Servicing. Their framework contains a number of assumptions and simplifications. They include uncertain customer revenue, but they do not include uncertainty for future launch costs or account for technology obsolescence. Established in the early 2000s, the framework does not include novel CONOPs for LEO-based OOS, such as temporary abandonment/orbital storage, pro-

liferated services, or combined OOS/orbital spare strategies. Additionally, they do not consider the combinatorial effect of flexible options.

Flexibility could benefit the OOS infrastructure as well as the satellite constellations it serves, since demand is highly sensitive to uncertain parameters such as launch and spacecraft costs. As de Weck et al. demonstrate in their flexibility framework of the 2000s Iridium constellation, incremental deployment that responds to the present state of demand could have reduced Iridium's losses by 30% [11]. Therefore, flexibility not only drives customer demand for OOS, it could improve the decision-making and design of the OOS infrastructure itself. It is therefore critical to consider both advantages to properly assess the value of flexible options and strategies.

No existing flexibility framework for LEO-based OOS includes multi-domain uncertainty, incremental decision making, and the combinatorial effects of classic OOS options as well as novel options, such as temporary abandonment, that would allow the user to screen for economically and environmentally feasible strategies and policies. However, frameworks for other infrastructures do address these needs and could be adapted. A close analog is a framework for offshore oil rigs that features a screening methodology that uses a mid-fidelity simulation, Monte Carlo scenarios, and decision rules to compare the value of offshore oil infrastructure options [12].

This thesis makes two primary contributions to advance sustainable space operations in Low Earth Orbit. First, it introduces Collection-as-a-Service (CAAS), a novel operational concept that transforms traditional multi-echelon sparing strategies into flexible collection hubs capable of satellite refurbishment, active debris removal deployment, and incremental capability upgrades. CAAS addresses the fundamental economic barriers to LEO-based on-orbit servicing by building upon proposed sparing strategies that aim to appeal to constellation operators, creating immediate operational value while preserving pathways for future technology adoption. Second, the thesis develops a comprehensive flexibility framework that integrates a discrete event

simulation with parametric policy modeling and adaptive decision rules to systematically evaluate technical configurations, flexible deployment strategies, and policy interventions across deeply uncertain futures. This framework represents a methodological advancement beyond previous OOS analyses by capturing multi-perspective decision-making, multi-domain uncertainty evolution, and the combinatorial effects of flexible options, which are essential capabilities for realistic infrastructure investment analysis but absent from existing frameworks. The analytical framework itself constitutes a standalone contribution applicable beyond the specific CAAS concept and use case used within this thesis. Structured around discrete event simulation with object-oriented architecture, the framework implements three concurrent decision-making scales: satellite-level decisions (repair, refuel, deorbit choices), fleet-level decisions (constellation-wide capability upgrades), and infrastructure-level decisions (warehouse additions and capability enhancements). Through systematic Monte Carlo experimentation across 80 uncertainty scenarios spanning technology costs, failure rates, policy environments, and market conditions, the framework employs Value at Risk/Gain (VARG) analysis combined with ranking convergence assessment and statistical significance testing to distinguish between genuine performance differences and statistical noise.

To demonstrate the framework's capability, this thesis applies it to a OneWeb-inspired mega-constellation case study in LEO. The case study examines an 18-plane, 648-satellite constellation at 1200 km altitude with 86.4-degree inclination, evaluating 29-43 unique configurations (depending on the experiment) across varying warehouse counts, satellite upgrade pathways (refuelability, repairability, RPO capability), infrastructure flexibility mechanisms, and eight distinct policy schemes over 30-year operational timelines. While this case study provides concrete validation of the framework's analytical capabilities and yields specific insights for OneWeb-class constellations, the framework itself is designed for broader applicability. Future ap-

plications could examine alternative constellation architectures (different altitudes, inclinations, or satellite counts), explore technology-specific servicing concepts (such as in-space manufacturing or modular satellite designs), or evaluate entirely different policy environments. The OneWeb case study serves to establish proof-of-concept for the integrated methodology while demonstrating how the framework illuminates the multi-dimensional trade-space where technical, operational, and policy decisions interact to shape outcomes for circular space economy development.

The experimental methodology progresses through three hypothesis-driven phases that systematically build understanding of CAAS performance. Experiment 1 establishes the baseline value proposition by comparing CAAS configurations against traditional overpopulation and multi-echelon sparing strategies, isolating the contribution of join active collection capability, ADR deployment, and Earth-return refurbishment. Results demonstrate that single-depot CAAS with RPO-capable satellites achieves -6.48% emissions reduction ($p < 0.001$) while maintaining cost parity with baseline (+0.76%, $p = 0.74$ not significant), validating that sustainability improvements are achievable without economic penalties. Experiment 2 introduces flexible deployment mechanisms such as adaptive satellite upgrade timing, conditional warehouse expansion, and incremental capability additions, to determine which flexible strategies improve performance under uncertainty. Ranking convergence analysis identifies flexible satellite serviceability upgrades ($R_f=1$) paired with pre-initialized satellite RPO capability as the consistently superior configuration, though absolute cost differences remain too small for statistical significance with 80 scenarios. Sensitivity testing confirms configuration robustness across operational cost perturbations and uncertain variable ranges, with satellite failure rates emerging as the critical determinant of CAAS competitiveness. Experiment 3 evaluates eight parametric policy schemes, from cost-neutral orbital use fees with collection rebates to growth-oriented subsidy mechanisms, identifying scenario-dependent optimal interventions. Cost-neutral Pol-

icy 1 (\$75k OUF with refund) achieves 29.47% refurbishment throughput increase ($p = 0.038$) without measurable cost impact, while Policy 1 applied to 4-depot infrastructure (\$50k OUF) delivers -9.94% emissions reduction at \$350M cost increase, quantifying the approximate business case gap that government subsidy or dual-use revenue must close to enable ambitious sustainability outcomes.

The integrated experimental findings validate that Collection as a Service represents a viable pathway toward circular space economies in LEO, though with important context dependencies and trade-offs. Single-depot configurations with flexible satellite upgrades and pre-initialized RPO capability emerge as the recommended architecture for cost-conscious stakeholders, achieving sustainability improvements through partial constellation servicing. Earth-based refurbishment consistently outperforms space-based servicing across both economic and environmental metrics, redirecting technology development priorities from complex on-orbit servicing toward reusable launch vehicle architectures and warehouse-based collection operations. Cost-neutral policy mechanisms demonstrate that sustainable infrastructure development need not impose permanent fiscal burdens, while sustainability-oriented policies reveal the \$350M investment threshold necessary for more ambitious 4-depot configurations that deliver -9.94% emissions reductions. Critically, the framework reveals that no single element independently guarantees success across all metrics; rather, its value lies in illuminating the multi-dimensional trade-space where technical configurations, flexible deployment strategies, and policy interventions interact to shape outcomes. These findings provide policymakers and industry stakeholders with quantitative evidence for context-dependent pathways toward circular space economies, demonstrating that economically viable sustainability improvements in LEO operations are achievable through strategic combinations of technical innovation, adaptive deployment, and targeted policy intervention, which ultimately advances the thesis objective of incentivizing private sector investment in space infrastructure that

balances economic viability, environmental responsibility, and operational resilience.

CHAPTER 1

INTRODUCTION

After decades of scientific and engineering breakthroughs, we've learned a great deal about our place in space. We've gazed into the universe, discovered new planets, built better communications, developed GPS, connected the world with internet, and gained better understanding of Earth's climate. We've also landed ourselves a new environmental hazard. Earth's orbits have become a tragedy of commons; providing benefit to many who bear no incentive to play caretaker. As the number of collisions in space has increased, the conversation about space stewardship has gained prominence. Despite the international conversation, however, there is little consensus on responsibility or the best path forward. To date, the international community hasn't agreed on a definition for space debris.

The current best practice is to deorbit satellites in Low Earth Orbit and push Geosynchronous satellites into a graveyard orbit. In recent years, the FCC has made moves to establish these best practices into their licensure requirements, imposing fines for those who do not comply with contractual end-of-life agreements [13]. A new issue arises from this solution, however. As mega-constellations come to term and operators begin deorbiting thousands of satellites on a continual basis, the population of pollutants in Earth's atmosphere will continue to grow. Single-use, short-lived satellites provide a flexible and cost-effective option for operators, but they contribute to a potential environmental issue. Like any conversation about sustainability, the true solution is to transition the linear economy of single use into a circular system that retrieves the full value of resources. A circular economy in space requires on-orbit servicing (OOS) capability. While OOS has had numerous successful demonstrations over the past decades, most private OOS ventures focus their attention on satellites in

geosynchronous orbit, which provides a more immediate business case than servicing mega-constellation satellites in LEO. There have been developments in OOS CONOPs for LEO, but these innovations struggle to compete with on-orbit sparing strategies that mitigate lost coverage due to satellite failures in a timely and inexpensive manner.

When it comes to building a business case for a circular system, there are numerous sources of uncertainty to consider. While uncertainty creates a challenge for architecture design, it also presents opportunity. Originally developed for quantitative finance, real options allow the decision-maker to take advantage of upswings and mitigate the effect of downswings by providing the right, but not the obligation, to exercise a particular option. Real options have since spread to engineering systems and evolved into frameworks that help engineers leverage flexible strategies and mechanisms to improve project outcomes. By anticipating future policy, future demand, and future impact to the upper atmosphere environment, satellite constellation operators may come to find that circular systems are better in the long run.

This thesis considers the role of flexibility and policy in OOS for LEO-based satellite constellations and the ADR vehicles that collect their failed satellites. Not only does flexibility improve the case for OOS because it provides value to customers, it could improve the manner in which OOS infrastructures are designed and deployed, taking advantage of opportunities and mitigating downsides as the future unfolds.

1.1 Thesis Objective & Contributions

The thesis objective is to identify opportunities that sufficiently incentivize the private sector to improve sustainability (via circularity) in LEO and reduce atmospheric emissions associated with single-use satellites. The three major contributions are a flexibility framework for LEO-based OOS options, novel OOS elements and CONOPs, and a methodology to assess policy impact on the OOS business case.

The flexibility framework to screen flexible options for LEO OOS infrastructure

leverages a discrete event simulation with object-oriented programming to manage the logistics, events, and interactions of the OOS infrastructure over each decision period. The simulation approach provides the necessary degree of fidelity while allowing the user to adequately compare the value of flexibility for various options. The framework includes multiple sources of uncertainty (launch cost, constellation growth, customer revenue, technology obsolescence, and random failures) and models demand as a function of uncertain variables and customer willingness to pay for servicing, which will be governed by customer decision making. The system of decision rules, embedded within the Discrete Event Simulation, allows for incremental changes in the infrastructure that respond to the present state of uncertainty and available capabilities in both satellites and OOS infrastructure. Monte Carlo scenarios, sampled from the multiple sources of uncertainty, will allow the user to conduct statistical analysis and compare the effectiveness of different options and strategies to identify which flexible options and strategies provide consistent benefit in managing total costs while providing environmental benefits like reducing total emissions, increasing servicing throughput, or reducing the number of satellites that deorbit after a single lifecycle.

The novel LEO-based OOS design element, collection hubs, allows for temporary abandonment in LEO, which is currently possible in GEO due to graveyard orbits but not feasible in LEO due to congestion and short orbital lifetimes. By collecting several old satellites in one location, the servicer may have the economies of scale necessary to justify servicing them. In the event that refueling and repairing these satellites is not cost effective, they can be safely deorbited. These collection hubs borrow and modify concepts from the Jakob et. al. multi-echelon sparing strategy, which contains a parking orbit for a spare warehouse and places orbital spares in each plane that are ready to respond quickly to random failures [14]. Multi-echelon sparing strategies leverage the timeliness of orbital spares while leveraging J2 effects to resupply these

orbital planes as the warehouse drifts across orbital planes. If active debris removal (ADR) vehicles collect old satellites and place them in the empty slots left in the spare warehouse as spares are deployed, these warehouses could refurbish old satellites and re-deploy them instead of launching batches of entirely new satellites. While a few researchers have speculated on the usefulness of junk collection in orbit [15] [16], none have fully investigated its feasibility or usefulness for creating an OOS business case.

Additionally, the framework also provides an opportunity to evaluate the effect of policy on demand for on-orbit servicing, so policy-minded users can understand how their policies might promote circular economies. Policy is an additional source of uncertainty for constellation operators, so including policy will allow the business-minded user to determine which infrastructure elements or strategies perform best under different policy types. While not enough is known about the environmental impact of atmospheric emissions, existing policy proposals for orbital congestion could simultaneously encourage the market for on-orbit servicing with little to no modification. This thesis will provide a parametric testbed for how different policies and policy parameters will improve the case for OOS in LEO. So far, research for this thesis has led to two related publications that were presented at Scitech 2025 [17] and Ascend 2025 [18]. The Scitech conference paper presented the novel CAAS CONOPs and preliminary flexibility framework while the Ascend conference paper leveraged the framework to conduct policy analysis. These conference papers are slated to become journal papers in 2026, following the completion of this dissertation.

1.2 Thesis Organization

Chapter 2 provides the motivation for this thesis as well as a literature review of relevant studies. It identifies the literature gaps, characterizes the problem, and presents the research questions along with the logic supporting the hypotheses. The logic flow is illustrated in Figure 1.3. Chapter 3 contains the methodology formulation,

depicted in Figure 1.3, necessary to address these hypotheses. Chapter 4 utilizes the methodology to address the research questions presented in Chapter 2, providing the experimentation results and analysis. Lastly, Chapter 5 summarizes the results, provides insights from the findings, reflects on lessons learned, and makes suggestions for future work within its concluding remarks.



Figure 1.1: Space Junk Allocation & Waste Abatement (JAWA)




Figure 1.2: Thesis Logic Diagram




Figure 1.3: Thesis Methodology

CHAPTER 2

BACKGROUND AND MOTIVATION

2.1 Growth of the Commercial Space Industry

Mega-constellations in LEO provide valuable services, such as Earth imaging for climate science and providing internet for under-served regions of the globe. Euroconsult forecasts that within the next decade, operators will launch 2,500 satellites per year, adding up to 24,500 satellites with a total market of \$400 billion [19]. A decade at this predicted growth means increasing the number of satellites by 4.5 times, increasing total mass in orbit by 3 times, and increasing the market value by 1.5 times. While constellation satellites make up 83% of the demand per satellite and will account for greater than 50% of additional satellites in the next decade, they only account for 30% of the market value.

Mega-constellations of small, inexpensive satellites in LEO are increasingly popular due, in large part, to decreasing launch and spacecraft costs. Proliferated constellations benefit from operational flexibility because they are able to reconfigure to mitigate the effect of uncertain demand and random failures. Including manufacturing and launch, GEO satellites cost between \$150 million and \$500 million and have an average lifespan of 15 to 20 years [20]. Meanwhile, the LEO heavy launch cost per kilogram has fallen more than 95% in recent years, from \$65,000 per kilogram to \$1,500 per kilogram—more than a 95 percent decrease [21]. Due to the growing trend in rocket reusability and heavy lift vehicles, Euroconsult reports that by 2031, they anticipate launch cost per satellite to be 33% of what it is today. LEO satellite cost and lifespan vary widely. The Starlink satellites, for instance, cost about \$500,000 and have a lifespan of 5 years. At full scale (12,000 satellites), Starlink would have a

reoccurring annual cost of approximately \$8.2 billion per year [22].

Observation 1: LEO Mega Constellation Growth

Low launch costs, the need for operational flexibility, and increased demand for space-based services are driving the number and size of LEO mega constellations.

The six major comsat mega-constellation companies, including Starlink, OneWeb, and GuoWang, contribute about 2/3 of the projected satellite total. Euroconsult also projects that the top six space-faring governments will still be the driving force in the satellite industry, accounting for 72% of the market value. A major risk facing the growing satellite industry is manufacturing overcapacity, since the combined capacity of the mega-satellite factories is approximately 10 times greater than Euroconsult's forecasted demand. Table 2.1 displays a summary of constellations projected for the next decade.

Table 2.1: Highlighted Constellation Projections

Constellation Owner	Number of Satellites	Altitude
Amazon	3,236	590-630km
Boeing	147	1056-44221 km
China SatNet	12,992	500-1145km
OneWeb	6,372	1200km
SpaceX	29,988	550km
Telesat	198	1000km

2.2 Orbital Congestion & Space Debris

Within the next few decades, at the current pace of space debris propagation, portions of space may be unusable. The US currently tracks 36,500 pieces of debris greater than 10cm. There are a predicted 500,000 pieces below 1cm. According to the ESA, the total mass of space objects in orbit weigh about 9,600 tonnes. The Starlink program alone will add 3,000 tonnes to LEO, which is greater than the current total

number of spacecraft in LEO. Starlink satellites have a 5-6 year life cycle and take 6 months to deorbit, which means that at any given time, 10% of the mega-constellation will occupy the same congested piece of space [23]. As more collisions occur, more collisions are likely to occur, leading to the phenomenon known as the Kessler Syndrome. A visualization of the crowded LEO regime is presented in Figure 2.2, where the blue dots represent active satellites, the red dots represent large pieces of debris, and gray dots represent small pieces of debris [24]. While mega-constellations did not create the entire space debris issue, they do contribute a considerable amount of mass to Low Earth Orbit and increase the risk of collision events.




Figure 2.1: Cumulative Objects in Space [23]

Observation 2: Orbital Congestion Challenge

LEO is a densely populated regime; mega constellations further exacerbate orbital congestion and collision risk.

2.3 Global Response to the Space Debris Problem

While many international leaders recognize the urgent need to rectify space debris and space congestion, there is little agreement over the proper course of corrective action. The space debris issue is in its adolescence, entering the mainstream conversa-




Figure 2.2: Visualization of Tracked Objects in LEO. Active satellites are blue, large debris is red, small debris is gray [24]

tions only 15 years ago. Addressing the space debris issue requires a reconfiguration of a deeply uncertain, international, and intricate socio-technical regime that has operated under the flag of exemptionism since its inception [25]. Space debris and space sustainability are often part of the same conversation, largely because sustainable practices could both mitigate and remediate space debris. Bringing sustainable practices to space requires a large inertial change, but industries and governments could apply lessons learned from other environmental efforts and energy transitions. True space sustainability will require international cooperation, new technologies, new business concepts, and new policies.

In the 15 years of international space debris conversation, values and perspectives have grown from mitigation (preventing more debris), to remediation (cleaning up current debris), and finally adaptation (finding uses for the debris) [26]. Yap et al. used socio-technical configuration analysis (STCA), a form of discourse networking, to track these trends. By categorizing discourse into state, market, global gover-

nance, global community, and sustainability values, Yap et al. tracked congruence between both concepts and actors. Starting in 2007 with the Chinese ASAT test, immediate attention was focused on the potential of space weaponization, shown in phase 1 of Figure 2.3. In 2008, a collision occurred when a defunct Russian satellite struck an Iridium satellite. In the following years, the conversation focused on the responsibilities of the state, disarmament, and methods to prevent further debris creation. The question of “who pays” gained central attention as spacecraft operators and governments considered the business case for debris removal. Consideration for space sustainability was not yet a central focus. During phase two, between 2012 and 2015, national interests took central focus as global community logic weakened, as shown in phase 2 of Figure 2.3. As governments and companies started to recognize the potential for space economies, they recognized that debris must be remediated instead of just mitigated. By 2016, more entities started discussing the need for sustainability both for and in space. This is clear in phase 3 of Figure 2.3, where we see sustainability logic take a central role in the conversation.

Over the last few years, a number of international organizations, such as the Inter-Agency Space Debris Coordination Committee (IADC), the United Nations (UN), and the International Organization for Standardization (ISO) have established guidelines for space debris mitigation. In 2022, the White House released its own Orbital Debris Implementation Plan, which describes specific actions in three categories: debris mitigation, tracking and characterization of debris, and remediation of debris [27]. Aligned with the findings of Yap’s STCA, the report notes its priority to improve safety and sustainability in the space environment. The inter-agency plan also includes intentions for research programs and development of debris recycling and re-purposing.

In 2023, the Space Safety Coalition released the Best Practices for the Sustainability of Space Operations report, which provides recommendations for best practices to




Figure 2.3: 3 Phases of the Yap et al. Socio-Technical Configuration Analysis show Sustainability Logic (purple) shift to a central focus by 2016 (Phase 3) [26]

address the growing needs of both space safety and space sustainability [28]. Proper end-of-life disposal is critical for debris mitigation; for LEO, this means deorbiting satellites within 5 years and for GEO, it means placing the satellite into a graveyard orbit.

Major US spacecraft operators have displayed commitment to adhering to the best practices. Although SpaceX hasn't signed the Space Safety Coalitions's report,

they have made efforts to properly deorbit their Starlink satellites, include collision avoidance systems, and reduce reflectivity of their satellites. Both government and industry are hopeful that “recommendations” and best practices are sufficient to properly mitigate space debris, but international adherence proves to be a struggle.

In 2022, the FCC turned best practices into rules, requiring all non-GEO spacecraft operators to reenter within 5 years of the mission period, a change from the historic practice of deorbiting within 25 years [2]. Additionally, the FCC requires spacecraft operators to keep impact energy below 15J and keep casualty risk less than 1 in 10,000 [29]. Practices to improve satellite demiseability decrease strike risks, but they introduce particles into the atmosphere [30].

The FCC has since demonstrated their commitment to these end-of-life disposal requirements. On October 2nd, 2023, the Federal Communications Commission fined Dish \$150,000 for failing to properly dispose of their EchoStar-7 [13]. Dish only made it to 122 kilometers above its geostationary orbit before running out of propellant, just halfway to their designated graveyard orbit. In 2012, Dish agreed to move its satellite 300km above geostationary orbit. In May 2022, they began to reposition EchoStar-7, not realizing that the spacecraft had less reserved propellant than expected. Although \$150,000 is relatively little, it marks the first time that the FCC has imposed a fine for failing to meet contractual end-of-life conditions. In addition to paying the fine, Dish must develop and install new orbital debris mitigation measures, such as better methods of monitoring propellant supply, more transparency about its disposal logistics, and better training for its employees. Loyaan Egal, the FCC Enforcement Bureau Chief, claimed that this penalty demonstrates how “the FCC has strong enforcement authority and capability to enforce its vitally important space debris rules” [13].

In a recent NASA analysis of space debris remediation strategies, key methods include ground-based and space-based lasers, as well as sweepers for small debris [31].

For trackable debris, options include controlled/uncontrolled reentry, rocket/laser just-in-time collision avoidance, and recycling debris for propellant. Major findings indicate lasers and just-in-time collision avoidance are cost-effective for small and large debris, respectively. Their study found that reusable remediation services for reentry show promise, while recycling debris into propellant for deorbiting operations lacks clear risk advantage. Notably, all methods rely on deorbiting for disposal.

Observation 3: Current Disposal Practice

Deorbiting is the only universally-accepted best practice to dispose of satellites in LEO.

2.4 Environmental Impact

Since space debris is a highly complex and interdisciplinary issue, John Pelton argues that a complete solution requires improvements in the following categories: technology, strategic management and planning, risk minimization, finance and capital planning, international policy and relations, and environmental assessment [32]. Environmental assessments for satellite constellations are usually focused on the space environment, but few consider the impact that these constellations of single-use satellites have on the atmospheric environment.

When it comes to space debris mitigation in LEO, the best practices are to employ collision avoidance systems on active satellites and deorbit end-of-life satellites. However, as mega-constellations grow, deorbiting satellites on the order of thousands could lead to several issues, such as increased strike risk, degradation to Earth's atmosphere, wasted terrestrial resources, and congested deorbit space. The NASA Space Debris Remediation Cost and Benefits study suggests that reentering space debris can release harmful materials and can catalyze harmful chemical reactions as the atmosphere gets hotter [31]. Although sustainable practices fell outside of the NASA report's scope, the authors did speculate that practices such as recycling debris could

offer environmental benefits, if not financial ones.

The Aerospace Corporation recently published about life cycle assessments of the space economy [25]. Jones and Jain suggest that, as the number of reentering satellites increase, particulates deposited upon reentry could change atmospheric behavior. As these particulates travel to lower altitudes, particularly alumina, they could “interact with ozone depleting chemistry” and contribute to changes in global temperatures [25].

Much like how the “Big Sky” theory was used to diminish concerns about crowding in space, the meteoroid theory addresses concerns about alumina particulates in the upper atmosphere. Given that alumina from meteoroids have had no measurable impact on Earth’s atmosphere, spacecraft operators and policymakers determined that deposited material from deorbiting spacecraft would be negligible. Boley and Byers challenge this existing environmental assessment and finds that alone, Starlink satellites will deposit more aluminum into Earth’s upper atmosphere than meteoroids [23]. On a daily basis, 52 tonnes of meteoroids enter the atmosphere, less than 1% of which is aluminum by mass. A 5-year cycle of a 12,000 satellite Starlink constellation would deposit 2 tonnes of aluminum per day [23], which would be roughly 4x the amount of daily meteoroid aluminum. In the past, scientists have suggested using aluminum at high altitudes in a controlled geo-engineering experiment to adjust the Earth’s albedo and counteract global warming [33], however, these suggestions were met with significant controversy [34]. Large-scale spacecraft deorbiting would effectively become an uncontrolled form of this controversial experiment [23].

Jain and Hastings consider a theoretical scenario of satellite reentry over the South Pacific and employ the Whole Atmosphere Community Climate Model (WACCM) to estimate the direct radiative forcing of the deorbiting alumina particles [30]. They assume that 1.33 Tg of alumina reenters the atmosphere on a yearly basis, matching the expectations for deorbiting mega-constellations. They represent alumina particles

as terrestrial dust in WACCM, since the software doesn't include alumina materials. They note that the current alumina flux levels are 150 Mg/year, which comes from assumptions about aluminum composition in satellites and rockets and how much mass ablates upon reentry. They also assume that all aluminum particles oxidize to form alumina, which is likely an overestimation. Another conservative assumption is that alumina particles form between 0.01 and 0.1 microns; conservative because large particles take the longest to descend. They repeat the aluminum flux over several years to resemble repetitive nature of reentering mega-constellations. Considering radiative forcing of alumina particles only, their findings suggests that there is no need to drastically alter best practices for satellite decommissioning because they find no long term issues related to reentering alumina's directive radiative cooling. However, they note the possibility for secondary consequences like the impact of asymmetrical radiative forcing on climate. They also note that there could be unintended consequences related to stratospheric ozone depletion. They don't consider other forms of atmospheric emissions related to deorbiting satellites and rocket launches, such as NO_x , black soot, hydrochlorocarbons, etc. They don't consider the emissions related to the rocket launches necessary to deliver each iteration of single use satellites. Therefore, due to its limited scope, this study does not serve as comprehensive evidence that deorbiting single-use satellites is acceptable in the long term and at scale.

Another study, entitled "Future Decreases in Thermospheric Neutral Density in Low Earth Orbit due to Carbon Dioxide Emissions", studies the impact that increased CO₂ concentrations have on neutral thermospheric density, which in turn would increase the amount of time needed to deorbit spacecraft thanks to less atmospheric drag [35]. The study employed a simulation of Earth's atmosphere and found that there's already been a 17% reduction in Earth's atmospheric density at 400km. At their 50:50 probability threshold, this number could reach 30% reduction if global warming stays within the bounds of the Paris Agreement (1.5 Celsius increase). This

corresponds to LEO orbital lifetimes that are 30% longer than they would have been in 2000, which increases the probability of collision and exponential debris growth. To meet the deorbit guidelines, spacecraft operators who rely on atmospheric drag will have to lower their final orbit perigee or implement new technology to increase drag, such as drag sails. Both of these options require more spacecraft mass to be dedicated to deorbiting operations.

Observation 4: Environmental Impact

Launches and deorbiting contribute to atmospheric emissions and environmental degradation; the extent of which is not entirely understood and largely unregulated.

Both rocket launches and reentering spacecraft create emissions that are harmful to the ozone and cause radiative forcing [36]. All combustion systems emit H_2O and NO_x while some also emit black carbon, alumina particles (Al_2O_3), and chlorine gas. Reentering spacecraft and debris also released thermal NO_x into the mesosphere. These emissions are detrimental to the ozone, as Ryan et. al explain in their journal paper [36]. The increased activity in the space sector has lead to notable changes in the stratospheric ozone layer. The black carbon (BC) released during rocket launches poses a greater threat to global warming than other forms of soot. As the space industry grows, it could undermine all the progress made by the Montreal Protocol [36].

In order to calculate pollutant emissions, Ryan et. al use activity factors and report emission pollutant factor by fuel type, assuming that all fuel is used at each stage. For reentering reusable vehicles, they estimated that NO_x emissions are roughly 17.5% of its mass, which is consistent with Larson et. al [37] and studies on Space Shuttle reentries [38]. For both controlled and unplanned payload/debris reentries, they estimate that NO_x emissions are equivalent to 100% of its mass. Their model found

that ozone in the upper stratosphere, at northern altitude, was the most susceptible to launch and reentry emissions. In the last decade, the biggest contributors to this depletion is 51% due to NO_x (reentry heating), and 49% due to chlorine from solid fuel rocket. Regarding reentering spacecraft, their biggest sources of uncertainty were location and mass of reentering objects as well as proper parameterization of NO_x emissions from reusable rocket stages [36]. This thesis develops an emissions model of NOx based on mass reentering the atmosphere.

Observation 5: NOx Emissions

While there are a number of different emissions related to rocket launch and atmospheric reentry, NOx is a strong proxy for this thesis because its formation and impact on the ozone are relatively well understood. The production of NOx emissions is related to mass reentering the upper atmosphere and vehicle demisability.

In short: deorbiting satellites have an effect on the environment, and changes in the environment have an effect on deorbiting satellites. While the impact of reentering satellites on the upper atmosphere is not yet entirely understood, it is evident that environmental considerations are worthy of consideration for future-minded space-craft operators. The National Environmental Policy Act (NEPA) applies to rocket launches, reentry, and recovery [39]. While regulations for environmental assessments of rocket launches are well established, there is less precedent for satellites in the space environment. To date, Earth's orbits are not yet considered a "human environment". The deorbiting strategy is sensitive to environmental and political changes, so it is valuable for constellation operators to consider system flexibility to account for the uncertainty.

One option to mitigate the environmental cost of LEO mega constellations without harming the growth of the space industry is to reduce reliance on single-use satellites

and extend satellite lifetimes, fostering a circular economy in space that makes use of the resources already in orbit and limits the need for deorbiting and launching new satellites. Not only could circular economies reduce the mass flux through the atmosphere, it could improve the bottom line for commercial space companies in the long run. Another option is to impose policies and regulations that specifically address atmospheric emissions and force private companies to modify their operations accordingly. The following sections discuss these two options and how they relate to the issues of orbital congestion and atmospheric emissions.

2.5 Circular Space Economies

Imagine discarding your car after it runs out of gas, or throwing away your laptop after it runs out of charge. Repeated use of expensive technology is not only cost effective, it's an effective use of Earth's resources. While the linear economy approach disposes waste after a single use, circular economies find a way to effectively recycle or reuse waste and get the most value out of resources. As the EPA defines it, "A circular economy uses a systems-focused approach that is restorative or regenerative by design to enable resources to maintain their highest value for as long as possible" [25]. Circular economies can be difficult to implement, but companies like IKEA, Burger King, and Adidas have made moves to improve the circularity of their supply chains, placing emphasis on reusing and recycling resources [40]. In some cases, policies are necessary to encourage private companies to embrace circularity, as we see with the lithium ion battery industry [41].

Circular systems are especially useful in applications with scarce and finite resources, which is the case for many space missions. For instance, the International Space Station is able to recover 98% of its water, recycling sweat and urine from the astronauts [42]. Circular Environmental Control and Life Support Systems (ECLSS) are of vital importance for extended, manned missions because they reduce the need

for resupply launches, reducing cost and extending mission timelines.



Figure 2.4: Elements of a Circular Space Economy

Just as circular economies improve ECLSS, they could both increase value for spacecraft operators and reduce the amount of mass flux through the atmosphere. While the exact impact of atmospheric emissions is unknown, reducing mass flux mitigates the uncertainty of atmospheric emissions. In doing so, the private company could improve its bottom line over the long run. As more is known about atmospheric emissions, policies could penalize spacecraft operators that don't operate sustainably, so it's wise for operators to be readily adaptable to sustainable practices. Figure 2.4 illustrates the elements of a circular space economy.

In order to reuse satellites, there must be capability for on-orbit servicing (OOS), which is a branch of on-orbit servicing and manufacturing (OSAM). As Jones et al. explain in “The Green Circularity: Life Cycle Assessments in the Space Industry,”




Figure 2.5: Circular Economies in Space [25]

circular space economies rely on the ability to service and recommission satellites at the end of their lifetimes [25]. They also highlight the value of reusing and recycling space debris left in graveyard orbits. Figure 2.5 contains their satellite lifecycle diagram.

Observation 6: Circular Economy Potential

Circular economies make efficient use of resources and reduce environmental impact, making them a key element for sustainable space industries. OOS is a critical capability for establishing circular economies in space.

Overall, the United States demonstrates sustained interest in On-Orbit Servicing

and Manufacturing (OSAM) development for economic, military, and sustainability purposes [43]. The 2021 Space Priorities Framework document notes that licensing requirements and operating guidelines would provide clarity and certainty to better support commercial OSAM [43]. In most policy documents regarding OSAM, however, US authorities focus on the need for more OSAM research and development to mitigate uncertainty. The following section reviews the state-of-the-art for OSAM and On-Orbit Servicing (OOS).

2.6 Current State of OSAM/OOS

The concept of OOS is nearly 20 years old and has been demonstrated in a number of government and commercial applications. In 2007, DARPA launched the Orbital Express mission, which successfully demonstrated the feasibility of autonomous operations to service a satellite [44]. Since the early 2000s, robotic OOS has been capable of the same repairs that astronauts made to the Hubble telescope [45].

In more recent years, the Space Force recognizes satellite servicing and on-orbit logistics as a “core capability”, planning to purchase in-space refueling and servicing as soon as the 2030s [46]. Prioritizing improved maneuverability and flexibility in the space domain, the Space Force is primarily focused on refueling services, but has also invested heavily in OSAM as well, expressing interest in propellant recycled from space debris. In February 2023, the U.S. Space Force awarded 1.7 million to Cislunar Industries, Astroscale, and Colorado State University for a Phase II Small Business Innovation Research (SBIR) award [47]. Ron Lopez of Astroscale U.S. said that the proposed CONOPs would provide a more “sustainable and flexible approach” compared to the old routine of launching and deorbiting assets.

Table 2.2: OOS/ADR Companies: Past, Present, and Future Missions with RPOD/RPO Capabilities

Company	Past Missions	Present/Ongoing	Future Missions	RPOD/RPO Use
Astroscale (JP/UK/US)	ELSA-d (2021-2024): First commercial debris docking demo, magnetic capture [48]	ADRAS-J (2024-present): Debris inspection of H-2A upper stage, achieved 15m approach [49]; ELSA-M: OneWeb satellite removal (2027) [50]	ADRAS-J2 (2027-2028): First large debris removal using robotic arm [51]; COSMIC (UK, 2025-2026): Remove 2 UK satellites [49]; LEXI GEO servicer (2025+) [52]	Autonomous RPO, magnetic docking, model matching navigation, angles-only navigation [53]
ClearSpace (CH)	None (founded 2018)	ClearSpace-1 (in development): Target changed from VESPA to PROBA-1 [54]	ClearSpace-1 (2026-2029): First ESA ADR mission, 4-armed capture, €86M [55]; CLEAR (UK, 2026+): Remove 2 UK satellites [56]; Intelsat GEO life extension (2026-2028) [57]	Autonomous rendezvous, robotic arm capture (4 arms), proximity ops [58]
Northrop Grumman SpaceLogistics (USA)	MEV-1 (2020): First commercial GEO docking with Intelsat 901, 5-yr life extension [59]; MEV-2 (2021): Docked with Intelsat 10-02 [60]	MEV-1 & MEV-2: Continue providing station-keeping services in GEO	MRV + 3 MEPs (2026 launch): Robotic servicing with NRL arms, customers: Intelsat (2), Optus (1) [61] [62]; DARPA RSGS: Inspection, repair, relocation [61]	Advanced RPO, autonomous docking, robotic manipulation, inspection [63]
Orbit Fab (USA/UK)	Tenzing (2021): First commercial fuel depot test [64]	RAFTI flight qualification (2024): First refueling port, TRL-8, \$30K price [52]; 100+ RAFTI units sold/in production [65]	Space Force Tetra-5 refueling (2025-2026): First operational hydrazine delivery in GEO; DIU RAPIDS fuel depot & shuttle (2025+) [66]; UKRefuel (2027): LEO refueling with ClearSpace/Astroscale [67]	GRIP docking mechanism, cooperative RPO with RAFTI interface, autonomous approach [68]
D-Orbit (IT)	Multiple ION orbital transfer vehicle missions	RISE mission development: ESA's €119M GEO servicing mission [69]	RISE (2028+): Commercial GEO satellite docking, maneuvering, and release for life extension [69]	Robotic docking systems, rendezvous in space [69]
Maxar (USA)	Development work on satellite servicing tech	NASA OSAM-1 (cancelled 2024): Was developing robotic refueling for Landsat 7 [70]	Robotic arms for various servicing missions [71]	Robotic satellite servicing, autonomous RPO [70]
Airbus (EU)	Development of robotic servicing systems	Partnership with Astroscale UK: 100+ docking plates ordered (2025) [72]; Developing modular servicing systems [71]	Robotic maintenance, refueling, life extension systems [71]	Robotic systems, multiple docking methods (robotic & magnetic) [72]
Telespazio/Thales (IT/FR)	Extensive ground operations experience	Italian IOS Demo Mission (PNRR funded): Ground segment development [73]; Space USB project: Standardized interfaces [74]	START servicing spacecraft (by 2026): Refueling, component replacement, orbit changes [75]; EROSS IOD: Robotic orbital support [74]	Robotic docking, rendezvous, standardized interfaces [74]
Momentus (USA)	Vigoride orbital transfer vehicles	Multiple Vigoride missions for last-mile delivery	In-space transportation and servicing expansion [76]	Proximity operations, orbital transfers

Table 2.2 summarizes past, present, and future OOS/ADR missions across various companies and highlights the growing technological maturity within the industry. Several critical technologies for on-orbit servicing, such as Orbit Fab’s RAFTI with a TRL of 8, have been successfully demonstrated or have upcoming demonstrations. SpaceLogistics has successfully demonstrated advanced RPO with autonomous docking for life extension purposes, while Astroscale has been making strides in RPOD with non-cooperative targets for ADR purposes.

2.6.1 OOS Business Case

While there are a number of technological, political, and financial challenges facing orbital congestion and atmospheric emissions, several companies have made big strides towards fostering the circular space economy by developing active debris removal and On-Orbit Servicing technology. ADR and OOS have become increasingly inter-related concepts in recent years, with many OOS and ADR companies either collaborating with each other or developing both operations in tandem [77].

While the looming question of “who pays” lingers, OOS and ADR companies largely rely on government investment. There are commercially funded missions, however, such as the SpaceLogistics MEV program. SpaceLogistics successfully docked their Mission Extension Vehicle (MEV-1) with Intelsat IS-901 in February 2020, with Intelsat paying \$13 million per year for life extension services [78]. Their MEV-2 followed with another successful docking in April 2021. Building on these successes, SpaceLogistics is developing the Mission Robotic Vehicle (MRV), slated for 2024, which will incorporate a 7-DOF robotic arm for inspection, relocation, active debris removal, and repair capabilities [62].

Astroscale is advancing ADR capabilities through missions like ELSA-d (launched March 2021), which successfully demonstrated magnetic docking and autonomous navigation, achieving rendezvous within 160m from 1,700km away despite thruster

anomalies preventing full mission completion [79]. Their upcoming COSMIC mission aims to capture two inoperable British satellites by 2026 [80]. Orbit Fab, founded in 2018, is building a propellant supply chain in space, having established the first fuel depot in 2021. Their engineers claim refueling stations could reduce debris removal costs by 80% and potentially reduce CO_2 emissions by 2.3 Gigatons annually [81] [82]. Their Podracer mission, planned for 2026, will demonstrate proximity operations and docking technology using their RAFTI and GRIP interfaces [83]. However, cost challenges remain significant. MAXAR’s OSAM-1 program was shuttered after NASA removed funding when costs approached \$2 billion, which is nearly triple the original \$626-753 million estimate [84]. A major cost driver was refueling satellites not designed to be refueled, though standardization of refuel ports across the satellite industry could improve the business case considerably.

Some researchers suggest that space debris itself contains value and could help encourage a circular economy in space [85] [86] [87] [88] [89]. Private companies are actively developing induction foundry technology that would make on-orbit recycling possible [90]. NASA is also investing in space debris recycling technology, providing \$750,000 to both Cislunar Industries and Yolo Robotics to finance the development of vacuum-rated electromagnetic induction furnaces that could transform space debris into useful stock and filaments [90].

Most OOS/OSAM concepts, including those that consider recycling debris, focus on GEO, since GEO spacecraft are usually more valuable and debris takes longer to naturally deorbit. Sears and Ho present an integrated model of a GEO On-Orbit Servicing infrastructure that studies the effectiveness of ISAM (In-Space Additive Manufacturing) and material recycling [86]. They found that these capabilities improve the infrastructure’s ability to service random failures and significantly improve required resupply launch mass. Another GEO OOS concept, called Space JANITOR, presents a multi-stage deployment plan that builds towards a fully-functional space

debris recycling and re-purposing infrastructure [89]. The Space JANITOR concept includes orbital cleaning drones (OCD) to bring debris to the modular Recycling Hub and Base (ReHAB).

In their framework to optimize many-to-many OOS operations in GEO, Sarton du Jonchay expands upon OOS options by including specific services such as inspection, refueling, station-keeping, repositioning, repair, and mechanism deployment while also incorporating the tools necessary to perform these services, such as refueling apparatuses, observation sensors, dexterous robotic arms, and capture mechanisms [20]. Sarton du Jonchay notes that service vehicles could include one or many of these capabilities.

Despite the large population of LEO satellites, issues with orbital congestion, and growing need to improve space sustainability, there are currently no plans for OOS in LEO aside from active debris removal and technology demonstrations that are conducted at lower altitudes to save cost [91]. GEO satellites are attractive for OOS because they are expensive to launch, have long life-cycles, and have long decay rates. Additionally, GEO satellites left in graveyard orbits still have the chance at a second life, effectively providing operators with the option to temporarily abandon assets. As the collection of valuable satellites in graveyard orbits grows, OOS providers have greater economies of scale. The business case for private OOS in GEO exists, as demonstrated by the MEV program. In LEO, however, satellite operators prefer to deorbit and follow the make-use-dispose cycle. Graveyard orbits and the option for temporary abandonment is not currently feasible in LEO.

In summary, the growing mega-constellation industry is causing environmental stressors which could be mitigated with policy and/or establishing circular economies in LEO. OOS is a critical capability for system circularity, but there is little demand and no clear business case for OOS in LEO. Understanding the differences between OOS in LEO and GEO is critical to understanding what measures would incentivize

OOS in LEO. The first step is to modify OOS infrastructure design to meet the needs of the LEO regime. It is necessary to ensure that the infrastructure reduces reliance on deorbiting to the extent that its own emissions contributions are justified. To improve the case of OOS in LEO, it may be useful to leverage the value of flexibility and adapt orbital congestion policies to encourage sustainable practices in LEO.

Observation 7: Technology Readiness

Much of the necessary technology for OOS has reached maturity. Currently, GEO is more attractive for OOS infrastructures than LEO due to the economic and operational challenges of conducting OOS in LEO. However, OOS in LEO would address the growing issues of orbital congestion and atmospheric degradation. Furthermore, circular systems in LEO could prove to be more profitable over the long run.

2.7 Developments in LEO-based OOS

Motivating Question 1

How can we mitigate the environmental cost of LEO mega constellations while incentivizing economically viable on-orbit servicing infrastructure that supports the growth of a circular space industry?

2.8 Hypothesis 1: Collection Hubs

2.8.1 LEO-Specific OOS Challenges

LEO is a different environment than GEO, so it only follows that OOS in LEO would differ from OOS in GEO as well. At low altitudes, J2 nodal precession due to Earth's oblateness causes RAAN drift, significantly altering orbital trajectories. Additionally, atmospheric drag has a more pronounced effect. Out-of-plane maneuvers are inefficient in LEO compared to GEO, which means LEO has a greater use for sched-

uled servicing than on-demand scheduling. Some researchers have dismissed OOS in LEO because maneuver requirements in LEO make multi-plane servicing infeasible or because LEO satellites have too little value to be worth servicing [92] [93].

Observation 8: LEO OOS Limitations

There are unique challenges for OOS in LEO compared to GEO: RAAN/inclination changes are very inefficient, constellations are proliferated, and LEO satellites are relatively cheap, making them less attractive servicing candidates.

2.8.2 Improving OOS CONOPs for LEO

Motivating Question 2

How to design OOS CONOPs specifically for LEO?

Proliferated Services

Luu et al. assess the value of OOS for mega-constellations in LEO by modeling several scenarios within a utility tradespace [3]. They find that in some cases, OOS compares favorably to spare satellites. This is especially true for high satellite failure rates. They consider three OOS CONOPS: traditional, depots, and pods. The traditional concept features a servicer that launches along with all its necessary resources. The depot concept contains depots positioned around the constellation that provide fuel or service to the mobile servicer. It provides a significant improvement over the traditional concept, but its design greatly depends on the configuration of the customers' constellation, creating a complex traveling salesman problem [3]. Lastly, the pods concept features proliferated modules that contain fuel and parts, located throughout the constellation. The servicer carries its own fuel supply and tools necessary to provide service. When a customer satellite needs service, the servicing vehicle captures the closest pod and then performs a rendezvous operation to intercept the customer. The pods concept, originally introduced by SpaceLogistics/Northrop

Grumman as their mission extension pods program, was originally designed for OOS in GEO [62]. Determining that the pods concept provides the clearest advantage for OOS in LEO, Luu et al. focus on this concept.




Figure 2.6: OOS Pod CONOPs [3]

Using multi-attribute utility theory (MAUT) to consider the various qualities that contribute to OOS performance, they include timeliness, number of servicing events, and mass delivered. Included in their trade-space is both chemical and electric propulsion, as well as varying altitude and satellite failure rate. They compare the utility of the pods concept with ground spares and orbital spares. In the scenario where all attributes are weighed equally, they find that OOS provides no benefit over the spare strategies. When timeliness receives the highest weighting factor, OOS with chemical propulsion provides better utility over spares when the number of spares is low, especially when relative to the number of satellites. In most cases, however, spares are still the dominating strategy. In scenario 3, when the number of servicing events receives priority, OOS provides much greater utility than alternative spares. These results suggest that OOS could be useful for providing a large number of inspections when fusing electric propulsion. Lastly, scenario 4, which focused on mass delivered, found that spares provide greater utility. When the failure rate is increasing from 30% to 45%, there was a greater advantage for OOS for scenarios 1 and 3.

Table 2.3: Tradespace Scenarios [3]

Scenario results for current paradigm of system reliability			Tradespace scenario results by increasing failure rate		
Current paradigm: failure rate = 0.3			New paradigm: accepting higher failure rate, $Fr = 0.45$		
Scenario	Description	Best option	Scenario	Description	Best option
1	Balanced	Spares	1a	Balanced	OOS
2	Timeliness	Mixed	3a	Quantity	OOS
3	Quantity	OOS			
4	Mass delivered	Spares			

Observation 9.1: Proliferated CONOPs

Proliferated, scheduled services improve feasibility of OOS in LEO compared to sparing strategies in some scenarios

Leveraging J2

The same work that considered proliferated pods to improve OOS CONOPs in LEO also includes orbital maneuvering that leverages J2 [91]. In their CONOPs, the servicer makes a RAAN change to maneuver between clients. Servicers could change planes by doing so at the poles, but this maneuver is inefficient, especially in LEO. Due to Earth's oblate spheroid shape, gravity is higher at the equator. As satellites in prograde orbit pass over the equator, they slow down slightly and drift towards the west. Leveraging J2 effects was first proposed by Legge in 2014 [94]. Luu et al. propose leveraging J2 to assist with RAAN changes by using a Hohmann transfer to reach a lower altitude and using nodal drift until the servicer satellite has reached the desired customer plane. Luu et al. determine that this strategy improves maneuvering within LEO.

In their 2022 conference paper, Geiman and other authors from Orbit Fab note that there are clusters of potential refuel customers in LEO, positioned on various planes and inclinations [95]. They advocate for a shuttle and depot concept that optimizes depot placement relative to these clusters. Like Luu et al., they suggest




Figure 2.7: Orbital Phasing [3]

using RAAN drift to deliver both depots and fuel shuttles across orbital planes in order to reduce delta-v needs. This maneuver is sufficient if the customer doesn't mind waiting longer for service.




Figure 2.8: Depot Drift Time vs. Orbit and Delta-V [95]

O'Leary et al. present a comprehensive refueling architecture for Active Debris Removal (ADR) vehicles centered around debris clusters in Low Earth Orbit [81]. The

study identifies three major delta-V clusters containing approximately 445 high-risk debris objects, representing over 50% of debris mass in LEO, and proposes a Depot-Shuttle refueling infrastructure positioned within these clusters to enable repeated debris removal missions. By leveraging J2 perturbation effects to optimize RAAN drift timing and minimizing delta-V expenditure through strategic depot placement, their analysis demonstrates that refueling can reduce total debris removal costs by up to 80% (from \$16.1B to \$3.3B) while reducing the number of required rocket launches. The architecture enables hybrid propulsion or chemical propulsion ADR vehicles to perform multiple debris removals throughout their 15-year operational lifetime by refueling after each mission, dramatically improving the economics of large-scale debris remediation compared to single-use ADR vehicles.

Observation 9.2: Leveraging J2 Perturbations

Leveraging J2 improves feasibility of OOS maneuvering in LEO

Cooperative Maneuvering

Ikeya and Ho compare fuel requirements for traditional, non-cooperative OOS with cooperative OOS for a multi-plane and multi-target case study in LEO [4]. In the cooperative OOS CONOPs, both servicer and customer spacecraft can maneuver towards each other. In semi-cooperative CONOPs, the servicer and customer perform either phasing or inclination change. Ikeya et al. derive an analytic relationship for the mass ratio (final servicer mass/initial target mass) that makes cooperative maneuvering more efficient than non-cooperative maneuvering. They assume all vehicles use chemical propulsion and they don't optimize refueling schedules. Varying target number and inclination, they assess the sensitivity of this derived condition. They determine that if the ratio of final servicer mass and initial customer mass is greater than the critical mass ratio, then cooperative maneuvering saves fuel. This is the case when the servicer is heavy compared to its customers.




Figure 2.9: Cooperative Manuevering in LEO [96]

Other researchers, such as Zhao et al., have also studied cooperative multi-customer refueling CONOPs for LEO [96]. Zhao et al. include J2 perturbations in their formulation and include constraints to account for surplus propellant on the customer satellites. They develop a mixed integer nonlinear programming (MINLP) model to represent their refueling scenario. To solve it, they proceed to develop a two-level hybrid optimization method that pairs a hybrid encoding genetic algorithm (HEGA) with linear relative dynamic equation to incorporate the J2 effects. Just like Ikeya et al., Zhao et al. find that cooperative maneuvering is most efficient when services are much heavier than the customers. Overall, they found that cooperative maneuvering offers propellant cost savings in most cases. They suggest considering many-to-many OOS CONOPs for future work.

Observation 9.3: Cooperative Maneuvering

Cooperative orbital maneuvers, for certain ratios of servicer/customer mass, reduces fuel requirements for one-to-many OOS in LEO when customers and servicers use chemical propulsion

Collection

GEO remains the more popular option for on-orbit servicing and ADR thanks to the scientific value and long life cycle of GEO-based satellites. Another component making GEO so attractive is its graveyard of defunct satellites. By taking the option to push end-of-life satellites into higher orbits, spacecraft operators have created future possible value for themselves because the asset is still recoverable. If this principle could be applied to LEO, could there be greater value for LEO-based on-orbit servicers? Obviously, graveyard orbits aren't an option for LEO, but what about "landfill" collection hubs? Literature on the logistics of collection-as-a-service (CAAS) is limited, but a few researchers have noted its potential.

Share My Space is a French SSA company that specializes in real-time traffic mapping and satellite logistics support. Their 2017 conference paper, entitled "Systematic space debris collection using Cubesat constellation," details an incremental method of using constellations of 6U ion-thruster cubesats to collect small space debris and bring it to temporary collection centers on orbit [16]. At these collection centers, debris could be recycled, refueled, repaired, or reused, depending on the state of the debris and the availability of OSAM technology. Eventually, this system could scale up to handle large debris. Lucken et. al claim that "once the debris are concentrated in just a few on-orbit storage sites, it will be much easier to deal with them on the long term, for instance by reprocessing them on a space station in LEO... or send them to the Moon to sustain future lunar economy" [16]. The paper cites the ISS ATV and Bigelow inflatable as examples of existing space structures that could serve

as a collection module. The paper speculates about the effectiveness of these collection centers, but doesn't elaborate on where they should be located, how many there should be, their capacity or their modularity.

In their journal paper about electrodynamic propulsion vehicles for LEO debris removal, Levin et. al note that collection could be an alternative to deorbiting, specifically in the crowded regions at high inclinations [15]. Levin et al. note that EDDE vehicles could corral a yearly 30 tons of space debris that, if recycled at an on-orbit facility, could help pay for their own removal. Collection could not only serve as a market mechanism, it could also help mitigate risks associated with reentry, especially for large objects [15].

Observation 9.4: Temporary Abandonment

Collection-as-a-service (CAAS) replicates GEO's option for temporary abandonment (via graveyard orbit) and could create the economies of scale necessary to incentivize on-orbit servicing. The collection hub concept provides value through flexibility because it retains satellite value and gives the collector options for future reuse or recycling.

Spare Strategies & OOS

Queuing theory and inventory management are critical space logistics subfields for both on-orbit servicing and spare replacement. Repairing satellites using spare parts and replacing satellites with spares requires many of the same tools and techniques. In their journal paper "Semi-Analytical Model for Design and Analysis of On-Orbit Servicing Architecture," Ho et al. combine queuing theory with inventory management techniques to develop a semi-analytical model that demonstrates the ability to assess the performance of an OOS system (that provides part replacement) without depending on computationally expensive simulations [97]. The authors note that modular and standardized satellites will give rise to improved OOS systems and bet-

ter space sustainability. They recommend using permanent and reusable servicers over disposal servicers with a pre-defined mission.

Observation 9.5: Space logistics

Borrowing concepts from logistics, queuing theory and inventory management could improve the design of space infrastructures

LEO CONOPs Summary & Gap

Proliferated CONOPs don't sufficiently motivate OOS in many cases, providing benefit only when failure rates are high and priority is placed on timely inspections. Cooperative maneuvering leads to lower propellant costs in cases where the servicer is much larger than the customer satellites, but researchers have only considered a one-to-many concept and have not included electric propulsion in their formulation. The majority of mega-constellations use electric propulsion, so this does not provide a clear conclusion. Changing orbital planes by leveraging J2 perturbations is helpful, but doesn't resolve all the difficulties of maneuvering in LEO. Collection hubs could provide sufficient economies of scale, but there is very little investigation into the concept. While some developments in LEO-based OOS are promising, they struggle to compete with the present business practices of launching cheap, short-lived satellites and deorbiting them as needed. Furthermore, satellite constellation operators are disinclined to make their satellites refuelable or repairable on Earth if there are no active service-providers in space. Likewise, would-be OOS providers are not inclined to launch infrastructure unless satellites are capable of being refueled or repaired in space.

Observation 10: Servicer-Customer Chicken-Egg Problem

The OOS market in LEO suffers from a coordination failure where satellite operators won't invest in serviceable satellites without available services, and service providers won't invest in infrastructure without serviceable satellites, preventing mutually beneficial OOS markets from emerging.

Motivating Question 3

How to better incentivize OOS in LEO?

2.8.3 Understanding the Needs of Mega-Constellations

For customers, OOS functions to mitigate random satellite failure, resupply constellations, and provide flexibility against uncertain demand. OOS can also allow spacecraft operators to reduce safety and redundancy requirements [3]. Flexible re-configuration allows spacecraft operators to mitigate the risks associated with random failures and uncertain demand [98] [99]. OOS could further enhance operators' ability to reconfigure their constellation and evolve with changing demand.

Mitigating the risks and consequences associated with satellite failure is a chief concern for satellite constellation operators. The FCC's \$150,000 fine against Dish for failing to properly dispose of its EchoStar-7 satellite set a precedent for penalizing satellite operators who violate their end-of-life licensure agreements [13]. Future penalties may far exceed this amount, especially as the orbital environment becomes increasingly congested. OneWeb has demonstrated its commitment to responsible operations through its partnership with Astroscale on the ELSA-M demonstration mission, which will remove a defunct OneWeb satellite from orbit [100]. Currently, ADR services to LEO are relatively expensive since they involve one-off missions to specific orbital locations. The ELSA-M mission to remove the OneWeb satellite will cost a total of \$48 million [101]. Achieving economies of scale for LEO ADR missions remains challenging due to the complexities of orbital maneuvering. Satellite constel-

lation operators stand to benefit from developing more cost-effective CONOPS for contracted ADR missions, particularly since they may face mandatory ADR requirements in the future, either through direct regulation or because accumulated orbital debris creates collision risks that prevent new satellite deployments. Kuiper has expressed concern about the stringency of current failure penalties, recently petitioning the FCC to reform or eliminate the 5-year deorbit requirement, calling it “artificial and rigid” [102] [103]. This suggests growing industry apprehension about compliance costs as nascent constellations scale up. Developing more cost-effective failure mitigation strategies would improve industry-wide compliance and orbital sustainability.

Observation 11: Satellite Constellation Operator Priorities

Satellite constellation operators are primarily concerned with consistent coverage and keeping costs low.

OOS and ADR are relatively new failure mitigation strategies for mega-constellation operators. If collecting and disposing of failed satellites is not required, satellite operators are inclined to rely on established methods to replenish their constellations and maintain consistent coverage. In their conference paper from 1999, Cornara et. al compare the usefulness of three spare strategies: launching replacements as needed, placing orbital spares, and placing spares in a parking orbit [104]. Many of these strategies are tested and well-understood. Iridium, for instance, already uses the orbital spare strategy.

Observation 12: Sparing Strategy Competition

OOS competes with other LEO constellation strategies such as overpopulation sparing, on-orbit spare depots, and flexible reconfiguration.

Recent developments further improve the spare strategy. In their paper, “Optimal satellite constellation spare strategy using multi-echelon inventory control with stochastic demand and lead times”, Jakob et al. implement a multi-echelon (s, Q) -type

inventory program to improve satellite constellation spare strategy [14].




Figure 2.10: Multi-Level Spare Strategy

Their architecture models the customer constellation after OneWeb and includes “warehouses” at lower parking orbits and “retailers” at each orbital plane of the mega-constellation it services. The warehouses leverage J2 effects to “drift” between orbital planes over time, delivering orbital spares to each plane as orbital spares maneuver in-plane to replace randomly failing satellites. Modeling satellite failures with a Poisson distribution, they use genetic algorithms to minimize maintenance cost of the mixed integer non-linear programming problem, accounting for performance requirements, orbital characteristics, and location policies. Their optimal solution reveals the cost efficiency of batch launches and demonstrates the flexibility of warehouses at multiple parking orbits.

Jakob et al. assume that in-plane spares are immediately available, which is true

for an overpopulated orbit, but not necessarily true if the satellite is located slightly off-plane to prevent collisions, in which case, the satellite would be available in about two days [14]. They disregard this delay in their model. They also assume that the closest parking spare warehouse delivers to the orbital plane in need and that launches only deliver resources to a single parking orbit. Lastly, they assume that there is only one active delivery at a time and that spares are launched in batches.

Making a number of assumptions for fixed simulation parameters, they assign the annual satellite holding cost to be 0.5 M US\$ per satellite per year. Using OneWeb satellites as their use case, they assume satellite dry mass to be 150 kg, satellite manufacturing cost to be 0.5 M US\$ per satellite, and launch capacity to be 34 satellites.




Figure 2.11: Inventory Spare Strategy [14]

While researchers have considered spare parts in the context of on-orbit servicing [97], no one has considered combining the multi-echelon spare strategy with the collection hub concept. The multi-level spare concept already includes a warehouse of spare satellites located at a parking orbit, but these satellites do not necessarily have to be brand new. If old satellites are returned to these collection hubs and refueled for future use, they could replenish constellation supply with less need to launch replacements from Earth. This novel concept would still provide the timely satellite replacement in the event of failure while providing all the long-term sustainability

benefits of on-orbit servicing. Furthermore, if these spare warehouses could dock and refuel ADR vehicles, there would be less need to launch one-off, single-use ADR vehicles directly from Earth. Docked on a parking orbit warehouse, they could respond to satellite failures as needed.

Motivating Question 4: Private Sector Incentivization

How can LEO-based OOS CONOPs be combined with strategic sparing CONOPs to motivate private satellite constellation operators?

2.8.4 Graveyard Orbits and the Value of Temporary Abandonment

The key to motivating private satellite constellation operators to embrace LEO-based On-Orbit Servicing (OOS) lies in the ability to seamlessly integrate with existing strategic sparing approaches while incrementally building future opportunities for OOS. Taking inspiration from GEO operations, where graveyard orbits preserve the recoverable value of temporarily abandoned assets rather than permanently disposing of them, LEO can adopt a similar philosophy through dual-purpose spare warehouses.

Bringing GEO satellites back from the dead has been demonstrated multiple times with Northrop Grumman's Mission Extension Vehicle program. MEV-1 performed the first-ever commercial spacecraft docking with Intelsat's IS-901 satellite in the graveyard orbit in 2020, then brought it back into operational GEO service for five years [105] [106]. In April 2025, MEV-1 completed its service mission and successfully undocked from IS-901, marking the first undocking between two commercial spacecraft in the GEO graveyard. MEV-2 also successfully docked with another Intelsat satellite, IS-10-02, in 2021 for a similar five-year life extension service. Other graveyard-salvage missions have also been proposed, such as the ESA contract with D-Orbit for their RISE mission [107]. Slated for 2028, RISE will demonstrate safe rendezvous and docking with geostationary satellites and conduct its demonstration phase by rising to the geostationary graveyard orbit to rendezvous with a client satel-

lite.

While graveyard orbits in LEO are neither safe nor practical, it could be possible to store satellites for extended periods of time onboard a spacecraft. The multi-echelon sparing strategy already establishes spare warehouses in parking orbits that leverage J2 perturbations to drift between orbital planes, creating an efficient resupply network for constellation operators. By extending these warehouses to serve as collection hubs, operators gain access to a middle-of-the-road option that mirrors the value preservation principles of GEO graveyard orbits. Just as GEO operators maintain future servicing missions to recover assets from graveyard orbits because they retain recoverable value, LEO collection hubs would concentrate temporarily retired satellites in specific orbital locations, creating economies of scale that reduce individual mission costs.

This dual-use approach capitalizes on spare warehouses' existing capability to house space assets for extended periods. Rather than immediately deorbiting failed or obsolete satellites, operators can defer disposal decisions while waiting to see if evolving business cases make servicing economically viable. The concentration of multiple satellites in a few orbital locations creates the critical mass necessary for cost-effective servicing operations, much like how spare parts warehouses achieve economies of scale through centralized inventory management.

For constellation operators, this approach reduces risk while maintaining operational flexibility. They can continue their sparing strategies while gaining optionality for future value recovery, therefore preserving the possibility that technological advances or changing market conditions could make satellite refurbishment, component harvesting, or other servicing activities profitable. This incremental approach to asset disposal aligns with private sector risk management practices while establishing the infrastructure foundation necessary for sustainable LEO operations.

Observation 13: Collection Hub Concept

A middle-of-the-road option between deorbiting and servicing in space could involve collection hubs that temporarily retain satellites before Earth return or refurbishment, creating economies of scale and reducing individual mission costs.

2.8.5 Enabling Technology for Earth Return

The fundamental advantage of LEO collection hubs lies in the operational flexibility they provide constellation operators through multiple disposal and value recovery pathways. By concentrating satellites in orbital warehouses, operators preserve decision-making optionality while market conditions and technology capabilities evolve. If refurbishment never proves economically viable, operators retain the straightforward option to deorbit the entire satellite collection in a controlled manner, achieving debris mitigation objectives without additional complexity. Alternatively, operators can incrementally upgrade warehouse capabilities to provide in-space servicing and directly recommission collected satellites in orbit as servicing technology matures. The most intriguing intermediate option involves returning collected satellites to Earth for ground-based refurbishment. While individual LEO satellites are too inexpensive and proliferated to justify dedicated return missions, the economics fundamentally change when satellites are concentrated in collection hubs. Emerging reusable second stage technologies under development by SpaceX and Stoke Space, along with reusable entry vehicles like Dragon capsules and Sierra Space's Dream Chaser, are creating new possibilities for bulk satellite return operations. The operational efficiency of this approach becomes particularly compelling when satellite return can be integrated with routine warehouse resupply missions, effectively utilizing the return leg of logistics flights that would otherwise carry minimal payload. Stoke Nova is actively developing a reusable second stage that leverages both its inher-

ent reusability advantages and its downmass capabilities [108]. In their independent investment guide, the company highlights how their reclosable payload bay design enables the capture and return of spacecraft to Earth. Stoke explicitly positions this capability as enabling the reuse and refurbishment of space assets, describing it as a critical technology for advancing more sustainable space operations.

This Earth-based refurbishment pathway could provide the critical missing link in resolving the chicken-and-egg problem that currently constrains the OOS business case. If satellites are routinely returned to Earth for refueling, component replacement, and payload upgrades, spacecraft manufacturers gain clear economic incentives to design satellites with modular architectures, standardized interfaces, and enhanced serviceability features. This design evolution would simultaneously enable more cost-effective ground-based refurbishment and lay the technological foundation for eventual space-based servicing operations. For instance, the Orbit Fab RAFTI refueling port, already at TRL 8, was intentionally designed to streamline on-Earth refueling as well as space-based refueling. Spacecraft designed for easy space-based refurbishment would also streamline Earth-based refurbishment.

During its operational lifetime, the Space Shuttle program returned several satellites to Earth for study, repair, or refurbishment for redeployment, as summarized in Table 2.4. Although Earth-return ended with the retirement of the Shuttle, these early missions demonstrated both the technical feasibility and the scientific value of recovering spacecraft. In the case of the refurbished and resold Westar 6 and Palapa B2 satellites, the insurance company covering the two spacecraft recovered \$50-60 million of its \$150 million loss [109], paying \$10.5 million to NASA and Hughes for their rescue and recovery [110]. Aside from reuse, returned satellites provide opportunities for forensic analysis, investigations of on-orbit degradation, and assessments of how service lives might be extended. In some cases, returned satellites could be salvaged for parts or materials.

With the private sector now driving launch costs far below those of the Shuttle era, and with companies developing fully reusable rockets that are expected to offer competitive downmass capabilities, Earth-based refurbishment has re-emerged as a plausible strategy for extending the life and utility of satellites in low Earth orbit. Furthermore, thanks to advancements in autonomous RPOD capabilities, retrieving satellites in space is far easier than it was during the space shuttle period.

Table 2.4: Spacecraft Retrieved by Space Shuttle and Returned to Earth

Spacecraft	Retrieval Mission	Details	Refurbishment & Reuse
Westar 6	STS-51-A (Discovery) Nov 16, 1984 [111]	Communications satellite stranded in wrong orbit due to Payload Assist Module (PAM) upper stage failure [111]	YES - Refurbished and relaunched as AsiaSat 1 by Hong Kong-based AsiaSat on April 7, 1990, aboard Chinese Long March 3 rocket [111]
Palapa B2	STS-51-A (Discovery) Nov 16, 1984 [111]	Indonesian communications satellite stranded in wrong orbit due to PAM-D upper stage failure [111]	YES - Refurbished and relaunched as Palapa B2R by Indonesia on April 13, 1990, aboard Delta rocket from Cape Canaveral [111]
LDEF (Long Duration Exposure Facility)	STS-32 (Columbia) Jan 12, 1990 [112]	Materials science platform with 57 experiments testing long-term space exposure effects on materials, components, and systems [113]	NO - Designed for reuse but never reflown. The structure itself was treated as an experiment and intensively studied. Data provided baseline for ISS design and future spacecraft [114]. LDEF was designed to be reusable but was never reflown [115].
EURECA (European Retrievable Carrier)	STS-57 (Endeavour) June 24, 1993 [116]	ESA microgravity research platform with 15 experiments in materials/life sciences, space physics, and astrophysics [117]	NO - Designed for reuse but never reflown. Was designed to fly five times with different experiments, but following flights were cancelled [118].
SFU (Space Flyer Unit)	STS-72 (Endeavour) Jan 13, 1996 [119]	Japanese reusable spacecraft with materials science and biological experiments [120]	NO - Refurbished for museum display only. Transported to Japan and refurbished for display at National Museum of Nature and Science in Tokyo [121]. Designed to be reusable to save money but never reflown [121]

The transition pathway from Earth-based to space-based servicing offers a risk-managed approach to circular space economy development. As satellites become increasingly designed for repair, featuring swappable payloads to address technology obsolescence and enhanced longevity, the business case for in-space servicing strengthens. Ground-based refurbishment operations would generate the demand signals, supply chain development, and technical expertise necessary to support eventual orbital

servicing infrastructure.

However, significant uncertainties remain regarding the costs, technical feasibility, and timeline for satellite return operations using reusable launch systems. The maturation of reusable second stage technology represents a notable source of uncertainty that could impact the viability of this approach. Nevertheless, the temporal characteristics of the collection hub concept may naturally align with these technology development timelines. Given the inherently slow pace of RAAN change via J2 perturbations, warehouses will require years to accumulate substantial satellite collections. This extended timeline provides opportunity for reusable return technologies to mature and demonstrate operational reliability before large-scale satellite return operations become necessary.

There are other challenges related to modeling downmass capability and availability, such as downmass payload capacity. While difficult to predict exactly, this framework assumes that download payload capacity is approximately one third of its upmass payload capacity based on SME insight. Another challenge of modeling reusable second stages is predicting what percentage of the launch vehicle fleet will feature reusable second stages and how the presence of a downmass business case may impact this percentage, since launch providers would then have a secondary revenue stream for their rockets and thus another reason to make their second stages reusable aside from their own operational cost benefits.

Several researchers are working to better understand the business case for reusable second stages and predict their adoption in the coming decades. In “An Economic Case for Distributed, On-Demand Orbital Down-Mass Systems,” Boysen et al. interviewed subject matter experts to forecast downmass services between 2030 and 2040 [122]. Their analysis identified large, low-frequency downmass systems as the primary market drivers, while also recognizing pharmaceutical, biotech, medical, and defense applications that could generate demand for smaller, on-demand, customized down-

mass solutions. The researchers noted that developments in space manufacturing supply chains could further expand these use cases.

Observation 14: Emerging Technologies for LEO Return

Returning satellites to Earth for scrap or future refurbishment is becoming increasingly possible due to progress in reusable second stages and reusable capsules (Dragon capsule, Sierra Dream Chaser, Stoke, Starship, etc.). LEO satellites are too cheap and too proliferated to justify individual return missions, but batch return could prove viable. However, the costs and timeline of returning satellites back to Earth on a reusable second stage are highly uncertain.

Although the study did not specifically identify returning old satellites to Earth as a market segment, it did recognize downmass applications related to space station component repair, failure analysis, and disposal operations. The authors also observed that the growing population of orbital transfer vehicles signals an emerging in-space logistics network with distributed transportation and multiple destinations, which could drive demand for smaller, on-demand return services. They argued that current downmass capabilities for such applications are inefficient because they are intended for human occupants, and that addressing these inefficiencies could significantly reduce costs for cargo-only operations. The study noted that Starship's anticipated full reusability will drive down return costs alongside launch costs, forcing smaller operators to both reduce their prices and differentiate through specialized niche services. Collection-as-a-Service (CAAS) exemplifies the type of innovative space logistics application that Boysen et al. predicted would drive demand for distributed, on-demand downmass services. Reusable second stage vehicles can efficiently combine upward logistics (resupply the orbital warehouses) with downward logistics (returning collected satellites) in the same mission cycle. This dual-purpose mission profile is an example

of a distributed in-space logistics network that maximizes operational efficiency while addressing the growing need for both orbital resupply and debris removal, contributing to the circular space economy.

2.8.6 Gaps and Proposed Solutions

The existing literature on LEO-based on-orbit servicing reveals two critical knowledge gaps that fundamentally limit our understanding of sustainable space operations. First, while theoretical proposals for collection-as-a-service concepts exist, no comprehensive feasibility analysis has examined the economic viability of collection hubs in LEO or systematically evaluated the operational options these infrastructures provide to constellation operators, such as docking ADR vehicles to retrieve failed satellites, collecting satellites to be brought back to Earth, or upgrading the warehouses to service satellites in space. The absence of a cost-benefit assessment leaves the business case for on-orbit collection hubs largely unsubstantiated.

The O’Leary et. al. refueling hub concept shares architectural similarities with the collection hub concept [81]. Both approaches utilize strategically placed orbital infrastructure leveraging J2 drift to enable efficient operations across multiple orbital planes and utilizing depots for multi-use ADR vehicles. However, while Orbit Fab’s architecture focuses exclusively on propellant resupply to extend ADR vehicle operational capability, CAAS provides a more comprehensive service suite including spare satellite storage, spare deployment, and potential refurbishment capabilities. The CAAS warehouses serve as temporary “abandonment” locations analogous to GEO graveyard orbits, collecting defunct satellites for potential reuse rather than immediate deorbiting, thereby creating additional economies of scale for on-orbit servicing operations. While the O’Leary et. al. refueling hub study focuses on remediation of existing debris, the collection hub concept is specifically focused on satellite constellation sustainability and creating circular space operations.

Research Gap 1: Collection Hub Feasibility

No existing literature provides feasibility analysis and cost/benefits assessment of collection hubs in LEO and the options it provides.

Second, despite growing recognition of the environmental impact of space activities, existing literature fails to quantify how LEO-based OOS operations compare to traditional deorbiting practices in terms of atmospheric emissions, representing a significant gap in our understanding of the true sustainability implications of different end-of-life satellite management approaches.

The CAAS infrastructure along with its future on-orbit servicing capabilities are intended to reduce atmospheric emissions by reducing the number of deorbiting satellites and reducing the system's reliance on ADR vehicles and spare satellites launched directly from Earth. In order to build and sustain this infrastructure, however, there will be launches and deorbiting second stages through the atmosphere; not all of which will be reusable. Therefore, it is important to consider the breakeven point where reductions in the number of deorbiting satellites offset the additional mass flux related to sustaining the OOS infrastructure.

Research Gap 2: Environmental Impact Assessment

Existing literature doesn't consider the impact of LEO-based OOS on atmospheric emissions compared to traditional deorbiting practices.

These knowledge gaps motivate two fundamental research questions that are essential for advancing the field of sustainable space operations. The first research question addresses the economic dimension: which collection hub infrastructure configuration provides the greatest economic value for OOS providers compared to traditional practices and other sparing strategies? This question requires systematic evaluation of various operational architectures, considering factors such as warehouse number, servicing capabilities, and satellite return options. The second research

question tackles the environmental dimension: which collection hub infrastructure configuration achieves the greatest reduction in atmospheric emissions and increase in satellite lifetimes compared to traditional overpopulation sparing strategies? This question requires emissions modeling that accounts for launch activities and end-of-life disposal across different operational scenarios.

Research Question 1: Economic Value

Which collection hub infrastructure configuration provides the greatest economic value for service providers and satellite constellation operators compared to traditional practices?

Research Question 2: Emission Reduction

Which collection hub infrastructure configuration provides the greatest reduction in atmospheric emissions compared to traditional overpopulation sparing strategies?

To address both research gaps and questions simultaneously, this research tests hypothesis 1, included below. This hypothesis directly addresses the first research gap by providing quantitative feasibility analysis of collection hub operations, while simultaneously tackling the second gap through comparative environmental impact assessment, considering sustainability metrics such as NOx emissions, number of de-orbited satellites, and number of refurbished satellites.

The research approach recognizes that economic viability and environmental sustainability are not necessarily exclusive objectives. This work seeks to identify operational paradigms that could incentivize private sector adoption while advancing environmental stewardship in space. The hypothesis specifically focuses on multi-echelon sparing enhanced with collection capabilities because this approach builds upon a sparing strategy intended to reduce costs for satellite operators, potentially reducing adoption barriers for the CAAS concept while providing the flexibility nec-

essary to adapt to uncertain future conditions in both technology development and market evolution.

Hypothesis 1

Multi-echelon sparing with collection warehouses that strategically retain salvageable assets for future reuse, refueling, or repair, while facilitating ADR missions and the controlled deorbiting or Earth-return of old satellites via reusable second stages provide better economic feasibility and sustainability metrics compared to infrastructures that rely on the overpopulation sparing strategy and ADR missions launched directly from Earth.




Figure 2.12: Illustration of the CAAS Warehouse Concept (*artwork by ChatGPT*)

Figure 2.12 illustrates the CAAS Warehouse concept, the details of which are explained in greater detail in Chapter 3. While the exact design and sizing of the spacecraft is left to future work, this thesis envisions a warehouse that resembles the Starlink payload structure, stowing satellites with their solar panels retracted. The spacecraft will have docking plates for ADR vehicles and allow for upgraded robotics

and refueling packages to enable in-space refueling and repair. The warehouse remains in orbit during a 30-year life cycle and is capable of passing its collection of satellites to a reusable reentry vehicle.

2.9 Hypothesis 2: Flexibility Frameworks

2.9.1 Flexibility Frameworks

The study of flexibility involves developing and studying novel designs and operations that could enhance an engineering system’s value by allowing operators to mitigate risks and take advantage of potential opportunities. Flexibility frameworks improve value as well as system sustainability and resilience [6]. While robust designs perform well under many circumstances without changes [123], flexible designs are designed to adapt to changes [124].

Since the early 2000s, researchers have been proposing flexibility frameworks for a wide variety of engineering systems and infrastructures. In general, they are composed of five phases: baseline design, uncertainty recognition, concept generation, design space exploration, and process management [6]. Flexibility is critical for sustainable design because it allows for the efficient use of resources in the face of high uncertainty. Compared to traditional design methods, flexible engineering design can improve economic and social performance by 10-30% or even more [6]. Rather than finding an optimal design for one future projection, it considers multiple future scenarios and improves the distribution of outcomes. Flexibility analysis measures flexibility value and compares the value of different flexible options. Typically, engineers are accustomed to equating uncertainty with risk – and when they do not include system flexibility, this equation holds [10]. Flexible designs, however, leverage uncertainty to improve system value and performance.

While the exact vocabulary differs in the literature, a flexible design question is twofold: how is the system designed to be flexible, and what design or operation must

exist to enable future flexibility? The former is called a strategy [6], a real option “on” a system [125], or a type [126]. The latter is called an enabler [6], a real option “in” a system [125], or a mechanism [126]. Some strategies are better at mitigating downsides, while others are better at taking advantage of upswings. The study of different strategies and enablers, along with their costs, forms the basis of a flexibility framework.

strategy	enabler
real option “on” a system	real option “in” a system
type	mechanism

The study of flexibility is rooted in Real Options Analysis, which quantifies flexibility value in large-scale, unalterable investment decisions. Real Options Analysis (ROA) got its start in finance, where Myers first introduced the idea as an investment decision making tool for real assets [127]. Myers’ Real Options concept adapted the Black-Scholes equation, provided below, which approximates the value of a financial option [128]. C is the call option price, S_t is the asset spot price, t is time to maturity, σ is volatility of the asset, K is strike price, and r_F is the risk-free rate. The Black-Scholes equation includes a number of assumptions to represent the call option cash flow as a stock purchased with borrowed money. Cox paired a simplified version of the Black-Scholes equation with binomial lattice analysis to provide a simple way of pricing financial options. [129].

$$C = N(d_1)S_t - N(d_2)Ke^{-r_F t} \quad (2.1)$$

$$d_1 = \frac{\ln \frac{S_t}{K} + t(r_F + \frac{\sigma^2}{2})}{\sigma\sqrt{t}} \quad (2.2)$$

$$d_2 = d_1 - \sigma\sqrt{t} \quad (2.3)$$

Several industries have adapted ROA for engineered systems, applying the technique to improve performance for expensive, long-term projects that experience a

high degree of uncertainty.

2.9.2 Real World Examples

Iridium

Several industries have adapted ROA for engineered systems, applying the technique to improve performance for expensive, long-term projects that experience a high degree of uncertainty. In some cases, ROA is applied retroactively to determine how a project might have gone better. Satellite communication company Iridium filed for bankruptcy in 1999, followed by Globalstar in 2002, thanks to the boom in terrestrial cellular networks, which exceeded 1991 expectations by 100% [11]. At the time that Iridium and Globalstar designed their large LEO communication satellite constellations, they expected a large demand for mobile satellite services (MSS) since terrestrial standards were lagging and cellular networks were making only modest gains. By the time these companies deployed their networks in 1998 and 2000, they realized their market predictions were greatly over-optimistic. Iridium faced greater than \$4 billion in debt while Globalstar faced \$3.34 billion in debt. While demand for MSS existed, it fell short of their constellations' capabilities. Failing to meet the expected demand, Iridium filed for bankruptcy in early 2000s, selling for less than 1% of their investment [6].

De Weck et. al studied the Iridium case and determined that incremental deployment of lower capacity satellites could have reduced their losses by 30% [11]. Rather than predicting global capacity demand, de Weck et al. introduce a flexible approach that deploys constellations gradually and with increasing capacity as demand increases. This approach reduces the possibility of economic losses associated with overly-optimistic projections. Additionally, they show that making constellations re-configurable also improves potential economic benefit, specifically in cases with high uncertainty. However, one detriment of flexible, incrementally-deployed

constellations is the associated cost of flexibility. Designers must be able to justify the cost of flexibility in their infrastructure design. Within this framework, the authors use Real Options Analysis as their analytical method to determine flexibility value. Incremental and re-configurable satellite deployment is an example of real options because it provides satellite operators with “the right, but not the obligation” to make a certain decision, such as launching more satellites, improving the capability of new satellites, moving existing satellites, or not doing anything at all.

As De Weck et al. demonstrate with their analysis on the Iridium constellation, flexibility frameworks are particularly useful for capital-intensive projects with a great deal of uncertainty. On-Orbit Servicing infrastructures certainly fit this description and stand to benefit from a flexibility mindset when it comes to their design and deployment.

Observation 15: OOS as Infrastructure Problem

OOS in LEO is an infrastructure problem with multiple sources of uncertainty including technology development, market demand, and satellite operator willingness to pay for service.

Flexibility in Infrastructure Projects

Across industries and applications, there are several examples of projects that have successfully implemented flexibility frameworks for their design and rollout. The 25 de Abril bridge in Lisbon, Portugal is an example of a successful flexible infrastructure. Originally, the bridge was built with the option to support 4 car lanes and a railway on the lower platform [130]. Today, the bridge has 6 car lanes and 2 rail tracks. This strategy deferred additional costs until there was a need for extra capability, leveraging the time value of money and lowering the Net Present Value (NPV).

The Health Case Services Corporation (HCSC) in Chicago, Illinois is another example of a flexible infrastructure. Originally built with 27 stories and the possibility

for future vertical expansion [130], they opted to complete the second phase when personnel needs increased faster than expected [6].




Figure 2.13: Flexibility Frameworks at Work

In “Enhancing the value of offshore developments with flexible subsea tiebacks,” Lin et al. consider three different flexibility types to address uncertainty in subsea tiebacks, finding that flexibility improves expected net present value by 76% [12]. They create a flexible option screening methodology capable of handling multi-domain uncertainty and a large combinatorial design space of strategies and design options by developing an integrated systems model and pairing it with Monte Carlo simulations and decision rules. They included multiple sources of uncertainty that were exogenous, endogenous, or a hybrid combination of the two. Lin et al. highlights the importance of including both types to fully understand the value of flexibility in capital-intensive projects [12]. The authors also note that interactions between different flexible options may be more valuable when others are already deployed. Evaluating multi-level flexibility is useful for large infrastructure projects because it

identifies advantageous combinations, interactions, and timing of options [12].

Studying water resources in Singapore, Zhang et al. employ integrated real options to evaluate state-of-the-art water management technology under uncertainty [131]. In “Analyzing Real Options and Flexibility in Engineering Systems Design Using Decision Rules and Deep Reinforcement Learning,” Caputo et al. analyze flexibility in a waste-to-energy system using deep reinforcement learning and discover a 69% increase in system value [132]. Cardin et al. determined that including incremental phasing, capacity expansion, and life extension in nuclear power plants amidst uncertain demand and public acceptance leads to better performing designs than those optimized with rigid strategies [133]. In a space logistics application, Chen et al. pairs network flow space logistics with decision rules in their multi-stage stochastic program to analyze ISS resupply and uncertain launch delays [134]. These framework examples are hardly exhaustive. Other applications include On-Shore LNG Production Design [135], on-demand vehicle-sharing, [136], power expansion generation planning [130], oil platform development [137], parking garages [138], and an IEEE 30 bus system [139], to name a few.

Observation 16: Flexibility Framework Success

Flexibility frameworks have been applied to analogous infrastructure projects and have identified options that bring additional value and reduce risk.

In these flexibility applications, both implemented and proposed, researchers have uncovered opportunity within uncertainty. Reducing infrastructure expenses and providing flexibility as a service could improve the feasibility of OOS from both the satellite-operator and service-provider perspectives.

Motivating Question 5

Could flexible options improve the demand for OOS in LEO?

2.9.3 Flexibility Frameworks for OOS

This thesis is not the first research effort to apply flexibility frameworks to OOS or consider the value of flexibility in the design of OOS systems. Recently, in their paper leveraging UMPIRE software to evaluate the value of refueling for several case studies, Burkhardt and fellow Orbit Fab authors argue that refueling satellites provides value via flexibility, allowing LEO constellation operators to re-configure to match evolving demand [140].

In the early 2000s, before the development and demonstration of OOS infrastructures, Saleh, Lamassoure, and Hastings authored a number of publications on flexibility in OOS [141] [142] [10] [9]. In their thesis, Saleh characterizes the nature, costs, and benefits of flexibility within space systems [143]. Elaborating on the importance of matching design lifetimes with the cycles of its dynamic environment, Saleh explains how OOS allows spacecraft operators to make decisions that better respond to an evolving environment. Later on, shortly after the Orbital Express mission, Nilchiani et al. proposed a framework that includes six fundamental elements to systematically analyze and measure flexibility within engineering space systems, using the Orbital Expression mission as their use case [144]. This framework adapts a similar set of assumptions and methods as those used by Saleh, Lamassoure, and Hastings. To the author's knowledge, these publications are the extent of academic work on OOS-specific flexibility frameworks.

In their series of papers and theses, Saleh, Lamassoure, and Hastings challenge the traditional approach of determining OOS feasibility, which was to compare the cost of servicing with the cost of launching a new spacecraft. As Saleh et al. point out, cost uncertainty obscures the truthfulness of these findings and the traditional approach neglects to consider the value of flexibility from the satellite-operator perspective [141]. The authors therefore propose a new approach that considers the value of servicing from the customer's perspective, essentially establishing the maximum

price that a customer would be willing to pay for servicing [141]. While focusing on customer perspective is useful for studying max price, it neglects to consider how flexibility could improve the rollout of the actual OOS infrastructure. This thesis seeks to consider both perspectives, making the simplifying assumption that the satellite constellation operator is also the owner/operator of the OOS servicing infrastructure that exclusively services its own constellation. This assumption is motivated by observed industry trends, as major constellation operators (e.g., SpaceX, OneWeb, Amazon) are increasingly vertically integrated, controlling launch, manufacturing, and ground infrastructure. Extending this integration to OOS is a natural progression and reflects emerging business models in the sector. Methodologically, single-entity ownership allows modeling of total system cost without the complexity of inter-firm pricing mechanisms, service contracts, or market equilibrium conditions between service provider and customer, while still modeling the decision rules governing flexible decisions in both the satellite fleet and servicing infrastructure. While future market evolution may see specialized third-party OOS providers, this assumption represents a plausible near-term scenario and establishes a baseline model. Extension to multi-party servicing markets with competitive dynamics represents valuable future research but would introduce additional layers of complexity beyond the scope of this thesis.

Over the course of two companion papers [141] [142], Saleh and Lamassoure develop a framework capable of identifying which missions benefit the most from OOS. They model non-deterministic random failures with failure rate λ , following a Markov process. To generate a cost model, they combine three standard cost estimation relationships: unmanned spacecraft, rule-of-thumb for industry, and small satellite. They perform a case study of the Iridium constellation in LEO, considering three different concepts: design 8-year-life-cycle satellites for replacement, design 8-year-life-cycle satellites for refuelling, and design 16-year-life-cycle year satellites with no need for replacement or refueling. In the event of random failures, the satellites are replaced.

To compare their flexibility framework with traditional methods, they develop a baseline model to evaluate the effectiveness of OOS. When the concept includes on-orbit servicing, the optimal baseline solution was to place two servicers per plane. Baseline results suggest that on-orbit refueling is only attractive under two conditions: (1) satellites operate at low altitudes where station-keeping demands substantial fuel, and (2) satellites have very low failure rates with fuel depletion as the dominant failure mode. Even in these cases, however, the cost uncertainty overshadows the benefit of refueling. Their flexibility framework improves upon this traditional approach by considering the flexible value that servicing provides.

Their framework builds on existing flexibility framework methods, such as decision tree analysis and real options theory to determine flexibility value. They include options to permanently abandon satellites, replace satellites, service satellites for life extension, upgrade satellites with new technology, or modify satellites to address shifting requirements or options. They assume a Markovian process for their uncertain parameter, which means history of the uncertain variable is not relevant. Saleh et al. argue that this is a reasonable assumption for exogenous uncertainty because it still reveals the impact of flexibility. They implement decision points within their framework where the formulation decides between different modes of operation.

They sum all mission and option costs into a cost metric, including discount rates for both costs and revenues that account for both certain and uncertain quantities. Certain quantities, or those attached to a twin security, are discounted with risk-free interest rate, r . Meanwhile, uncertain quantities, or amounts not attached to a twin security, have an internal rate of return that adds a risk premium to the risk-free interest rate. For the sake of convenience, they assume that the risk-free interest rate is constant throughout the mission timeline. They design their decision model to maximize future mission value.

They first consider the singular option of life extension, employing the Black-

Scholes equation. They quantify flexibility value as the difference between Expected Value and Traditional Value. They note that option value does not exceed the loss it potentially prevents.

When spacecraft operators choose to service a vehicle, they have enabled the option to service the vehicle again in the future, called a compound option. Their formulation includes a matrix of switching costs, dependent on time, that prevents impossible switches. For instance, abandoned assets can not be serviced again in the future, so its switching cost is set equal to infinity. At each decision point, the framework compares the cost and utility of each option and maximizes for expected value, taking the state of the uncertain parameter into account.

The authors make several assumptions, such as discretizing time into distinct steps to generalize continuous decision-making in the iterative backward process. They decide to make operational modes distinct, but this could be adjusted by using an integral to select from a continuous range of operational modes. Their framework applies for exogenous uncertainty, such as failures, market behavior, and developments in technology, which is not impacted by decisions-making.

However, OOS contains endogenous uncertainty as well, such as market dynamics and interactions between satellite-operator and service-provider. Customer demand for OOS is sensitive to both external factors and servicer decision-making. Meanwhile, servicers will make decisions about possible upgrades according to customer demand. These two agents experience uncertainty that is both dependent and independent of their decision-making, but this is not accounted for in these OOS frameworks.

Matos de Carvalho et al. recognize the need to capture client/servicer dynamics and develop an agent-based modeling and simulation (ABMS) framework to better understand how their relationship impacts OOS [145]. Their ABMS uses a set of rules, metrics, and characteristics to establish the relationship between servicer and customer. They create a hypothetical customer constellation and apply their frame-

work to explore OOS operations such as life extension, refueling, and rescue/recovery. Finding the client/servicer relationship impact to be nontrivial, Matos de Carvalho et al. highlight the importance of including multiple perspectives and multiple metrics. Since servicers are not inclined to make satellites serviceable in space if there are no available in-space services, and service-providers are not inclined to launch servicing infrastructure if satellites are not designed for in-space servicing, it is critical to capture the interaction between these two entities.

Observation 17: Satellite Upgrade Decisions

Making satellites refuelable or repairable is critical for OOS viability, but satellite operators are disinclined to do so unless they see immediate value-add. The timing and conditions for these upgrade decisions significantly impact OOS infrastructure rollout.

Regarding their OOS flexibility framework, Saleh and Lamassouire note that the framework is limited by the accuracy of the probability density functions it samples for its uncertain parameters. Lastly, the framework defines distinct types of flexible options but neglects to include the possibility of new modes in the future.

In their companion paper, they apply their flexibility framework to different mission types to evaluate the usefulness of flexibility in different applications [142]. Due to fuel requirements and maneuvering time, they find that servicing provides little benefit for a commercial radar constellation in LEO. For GEO communication satellites, however, they determine that flexibility helps improve capacity. The commercial LEO constellation case considers uncertain revenue and includes the options to abandon or service for life-extension. They represent uncertain revenues with a geometric walk model with drift and volatility parameters.

First, they consider the sole option of abandonment, assuming that abandonment costs operators nothing if deorbiting propellant is included in the spacecraft design.

In this initial study, they determine that traditional methods have underestimated mission value by failing to include the option to abandon, which increases as system uncertainty increases.




Figure 2.14: Value of Compound Abandonment [142]

Their second study includes the option to provide service for life extension. With OOS available, operators can choose to shorten satellite lifetimes to better match the length of market cycle periods [142]. Based on results from the abandonment only study, they determine that temporary and irrevocable abandonment have the same value. However, when the option to service is available, temporary abandonment has more value. Their framework neglects to consider the value of temporary abandonment when combined with other options. Additionally, these frameworks include OOS options available in the early 2000s, so they do not include recent developments and novel concepts for OOS in LEO. Including new options and combinatorial interactions may reveal greater flexibility value.

Predating these twin journal papers [141] [142], Lamassoure completed a PhD thesis on the same topic, evaluating the cost-effectiveness of on-orbit servicing and how the option to service spacecraft could reduce their designed life-cycle [10]. With a

similar flexibility framework, they conclude that when spacecraft operators consider the option for on-orbit servicing, the optimal satellite design point sees a shorter life cycle. For future work, they suggest including random failure and technology obsolescence.

In their thesis, Lamassoure explores various options for servicing infrastructures, including servicing/refueling depots that satellites either directly maneuver to for servicing or servicing vehicles maneuver to in order to refuel for their next servicing mission to a satellite that remains in place [10]. Lamassoure acknowledges the costs associated with changing orbital planes and concludes that servicing would be restricted to an orbital plane, not considering the effect of J2 perturbation. In the case of satellites maneuvering to the depot for service, they state that the down-time may be unacceptable. They do not consider possible coordination with multi-level sparing that would provide an immediate replacement and remediate the long down-time issue. Overall, they highlight the high risk of collision given the large number of RPOD maneuvers associated with the depot approach. However, over the last 24 years since this thesis was published, the TRL for RPOD has improved considerably.

Overall, flexibility frameworks prove to be useful for identifying opportunities that improve the design and deployment of On-Orbit Servicing infrastructures. OOS provides value to the customer by providing flexibility in their operations. Flexibility in the deployment of OOS infrastructure itself could provide a better bottom-line to the service provider, but this has yet to be explored. The existing OOS flexibility frameworks present a number of other gaps, such as including multi-domain uncertainty, considering recent developments in CONOPs, and incorporating the combinatorial effects of flexible options.

Research Gap 3: LEO Flexibility Frameworks

No flexibility frameworks exist for OOS in LEO that consider multiple sources of uncertainty, include CAAS/warehouses, or account for both exogenous and endogenous uncertainties, such as the interaction of decisions between satellite-operator and service-provider.

Research Gap 4: Incremental Deployment

Previous papers that propose LEO-based OOS concepts do not consider the merits of incremental deployment of OOS infrastructure that evolves with the uncertain environment or the consequences of delayed decision-making about satellite upgrades.

Research Gap 5: Combinatorial Flexibility Benefits

Existing OOS flexibility frameworks do not consider the combinatorial benefit of multiple flexible options.

2.9.4 Flexible OOS Concepts

Previous OOS Flexibility Frameworks identify a few flexible options, some of which are included in Table 2.5, but it is necessary to consider what additional options should be included and how they should be included [10]. Flexible design concepts are composed of two elements: strategies and enablers [6]. Common strategies, sometimes called real options “on” a system [125] or types [126], are abandonment, investment deferral, expanding or contracting system capacity, staged deployment, research and development investment, etc [6]. Enablers, sometimes called real options “in” a system [125] or mechanisms [126], represent the actual designs or operations that allow for flexible option implementation. Enablers require technical understanding of the system while strategies exist at the managerial level. There are several methods to determine which options to include in a framework, such as brainstorming, drawing from industry standards, using options screening, or using holistic methods [126].

Table 2.5: Space Mission Options [10]

Source of Uncertainty	Option for Traditional mission	Additional Option for Serviceable mission
Operations cost	Abandon	-
Random failures	Replace	Refuel, Repair
Market demand	Abandon-or-replace	Abandon-or-service
Technology	Upload software Replace	Upgrade (e.g. computer)
New requirements	New mission	Modify

This thesis will borrow from options included in previous OOS flexibility frameworks, explained in detail in subsection 2.9.3, and supplement them with recent OOS developments as well as novel concepts. This brings us to Research Questions 3:

Research Question 3: Optimal Flexibility Options

Which flexible option, or set of flexible options, provides the greatest economic value and environmental benefit for OOS providers under uncertainty?

The collection hub concept, introduced previously, is a flexibility enabler since it allows the service provider to build its supply of satellites over time and make incremental servicing decisions as the capability becomes available. Similarly, reusable second stage vehicles that can return satellites back to Earth are also flexibility enablers, since they extend servicing options and provide salvage value in the event the satellite is brought back to Earth but not recommissioned.

Observation 18: Warehouses and Reusable Second Stages as Flexibility Enablers

Collection warehouses are flexibility enablers that provide options for delayed upgrades and deployment, allowing adaptive responses to market conditions.

Combining multi-echelon spare strategies, cooperative maneuvering techniques, and collection-as-a-service, we get the following collection of CONOPs options, contained in Table 2.6 and Table 2.7. Bold cells represent objects or concepts that are explored in this thesis.

While there are a number of possible in-space services, like recycling and inspection, this thesis will focus on refuel and repair services since these are fundamental for a satellite's extended occupation in space. Repair services will include both failure repair as well as payload swap for obsolete satellites. Unlike previous OOS research, this thesis will include the option for satellite servicing to occur on Earth as well as in space. The framework will include the following flexible options:

1. Permanently Abandon Satellite (Deorbit)
2. Temporarily Abandon Satellite (Pay for Collection)
3. Receive service for satellite life extension (in space)
4. Send satellite to Earth
5. Receive service for satellite life extension (on Earth)
6. Deploy new collection hub/warehouse
7. Deploy servicing capability to existing collection hub

Previous on-orbit servicing (OOS) flexibility frameworks underestimate the value of temporary abandonment because they fail to account for its combinatorial benefit when integrated with other flexible options. Unlike traditional frameworks that treat temporary abandonment as equivalent to permanent abandonment, the CAAS concept leverages collection warehouses to enable true temporary abandonment, where satellites can be retrieved and potentially refurbished later when market conditions become favorable.

The incremental deployment of capability, both in terms of warehouse infrastructure and satellite capabilities, fundamentally improves the business case for OOS providers by creating a responsive, adaptive system. This approach enables satellite operators and service providers to react dynamically to each other's decisions and

Table 2.6: Morphological Matrix for OOS/ISAM Part 1: Architecture and Capabilities

OOS/ISAM Category	Options / Strategies				
ARCHITECTURAL ELEMENTS					
Regime	LEO	MEO	GEO	Multi-Orbit Parking	Multi-Regime
Concept	One-to-One	One-to-Many	Many-to-Many	Multi-Echelon	Proliferated Pods
Infrastructure Elements	Collection Warehouses	Mobile Servicers	ADR Vehicles	Reusable 2nd Stages	Service pods
Customer Elements	LEO Constellations	GEO Comsats	Gov't Satellites	Earth Observation	Scientific Sats
COMMODITIES & RESOURCES					
Warehouse Commodities	Collected Satellites	Bipropellant	Monoprop	EP Propellants (Kr, Xe)	Repair supplies
Servicing Tools	Refueling Apparatus	Observation Sensors	Robotic Arm	Capture Mechanism	
TECHNICAL CAPABILITIES					
Warehouse Capabilities	Old Satellite Storage	Spare Storage	CP Refueling	EP Refueling	Repair
ADR Vehicle Capabilities	Inspection	CP/EP Refueling	Station Keeping	Capture	Deorbiting
Service Location	In-Space Only	On-Earth Only	Space + Earth	Mobile/On-Demand	

evolving market conditions, fostering the customer-servicer interaction that previous frameworks have neglected. Rather than requiring large upfront capital investments in full servicing infrastructure, operators can gradually build capabilities as demand materializes and technology advances.

Strategic timing plays a critical role in risk mitigation. Providing OOS capabilities during a satellite's disposal period, when the satellite would otherwise be deorbited, creates a unique risk-benefit profile. If servicing attempts fail during this phase, operators lose only satellites they would have discarded anyway. This timing strategy transforms the traditional risk calculus of OOS investments.

While Hypothesis 1 seeks to quantify the value of CAAS relative to other spare

Table 2.7: Morphological Matrix for OOS/ISAM Part 2: Operations and Flexible Options

OOS/ISAM Category	Options / Strategies				
OPERATIONAL STRATEGIES					
Service Scheduling	On-Demand	Scheduled/ Periodic	Combination	Predictive (Condition)	Emergency Response
Decision Timing	Immediate	Delayed	Staged	Contingent	Adaptive/ Learning
Rendezvous Type ¹	Semi-Coop (Cust. moves)	Semi-Coop (Serv. moves)	Fully Cooperative (both move)	Non-Coop	
FLEXIBLE OPTIONS					
Option Type	Permanently Abandon (Deorbit)	Temporarily Abandon (Collect)	Service In-Space (Refu- el/Repair)	Return to Earth	Service on Earth
Upgrade/ Modify Options	Payload Swap	Subsystem Replace	Propulsion Add	Software Update	Structural Mod
Deployment Options	Deploy New Warehouse	Deploy New Servicer	Upgrade Existing Hub	Expand WH Capacity	Upgrade Satellites

deployment strategies, Hypothesis 2 addresses a different question: how does implementing flexible options enhance the overall feasibility of the CAAS concept? This hypothesis specifically focuses on identifying which individual options or option combinations provide consistent value across a wide range of uncertain futures. By examining the synergistic effects of multiple flexible options operating simultaneously, this analysis aims to determine whether flexibility can transform CAAS from a marginal spare strategy alternative into a compelling business case for sustainable space infrastructure.

Hypothesis 2

If the flexibility framework models CAAS system evolution, captures uncertainty, and models interactive decision-making, it will identify which flexible enablers and strategies for both servicing infrastructure and satellite constellation further improve the economic feasibility and sustainability metrics of the CAAS infrastructure compared to infrastructures that only rely on the overpopulation sparing strategy.

2.10 Hypothesis 3: Policy

2.10.1 The Policy Problem: Environmental Impact and Market Failures

The space industry's environmental footprint is uniquely complex, spanning terrestrial, orbital, and atmospheric domains. As launch rates accelerate and satellite constellations proliferate, understanding and mitigating these multi-domain impacts becomes increasingly critical.

While much of space sustainability research has focused on orbital congestion, a growing body of work now investigates atmospheric pollution from spacecraft reentry and burn-up [23] [36] [25]. This concern is compounded by the impact of climate change, which is altering upper atmospheric composition. This could potentially lengthen orbital lifetimes and complicate deorbit operations that rely on atmospheric drag [35].

Despite these growing concerns, atmospheric pollution has not yet constrained the growth or operations of satellite mega-constellations. Given the nascent state of research on the matter, private operators are not inclined to stray from the tried-and-true method of deploying cheap and short-lived proliferated networks of satellites. Creating regulations to address an environmental issue before it reaches mainstream attention is inherently difficult. Moreover, as Kuiper's recent Notice of Ex Parte filing to the FCC demonstrates, constellation operators show limited willingness to comply

even with existing orbital protection regulations [103]. To preempt environmental degradation, it is necessary to establish a compelling business case with sufficient financial benefits to prompt satellite operators to abandon their make-use-dispose model.

However, transitioning from linear to circular economies entails significant risk. Even if the CAAS concept paired with flexibility benefits proves cost-neutral or cost-advantageous, the private sector will likely require public backing to de-risk the transition. The Astroscale-OneWeb ADR contract exemplifies this dynamic: the mission is viable largely due to partial UK government funding [101]. Without demand signals from governments, sustainability improvements in space operations are unlikely to materialize organically. Yet subsidizing all sustainable practices would be both fiscally prohibitive and politically unpalatable. Effective government intervention must therefore be strategic, targeted, and aligned with constellation CONOPs to signal demand without imposing excessive costs or hindering industry growth.

Observation 19: Business Case Gap and Government Intervention

The private sector is unlikely to adopt sustainable space practices without government demand signals.

The following sections explore the current regulatory landscape, evaluate proposed policy mechanisms for curbing orbital congestion, and examine policy instruments from analogous industries to identify approaches that could facilitate more circular systems in LEO.

2.10.2 Current Policy Landscape and Regulatory Gaps

Although launch and orbital activities are subject to regulations such as the FAA's 1998 Commercial Space Launch Act Amendments and licensing requirements enforced by the FCC, current frameworks lack provisions addressing upper atmospheric pol-

lution linked to space operations [146] [2] [29]. Earth’s orbits are not considered a human environment, and the Montreal Protocol does not apply to the space sector, so the domain falls outside the jurisdiction of both the National Environmental Policy Act (NEPA) and the Montreal Protocol [36].

International Space Law and Debris Management

The 1967 Outer Space Treaty, signed ten years after Sputnik, established an agreement for the peaceful use of space and banned nuclear proliferation, but neglected to include a clear and enforceable definition of space debris [147]. Additionally, the Outer Space Treaty states that “registry State jurisdiction and control can only be transferred to another State, not to a private entity,” and that the spacecraft operator always retains ownership rights. This creates an obstacle for active debris removal (ADR), since it bars states from interfering with space objects registered to a different country.

At the international level, the United Nations Office for Outer Space Affairs (UNOOSA) and the United Nations Committee on the Peaceful Uses of Outer Space (COPUOS) are responsible for coordinating best practices and international agreements regarding the use of space. International consensus on space debris policy proves to be a lingering challenge. As governments implement their own policies, there is growing concern that private companies will go “forum shopping” if the regulations in their native country prove to be too stringent.

Attempts to Expand NEPA Jurisdiction

There have been recent efforts to include the space environment within NEPA jurisdiction, but to no avail. Arguing that Earth’s orbits should be considered as a human environment, Viasat, DISH, and The Balance Group submitted an Amicus Brief to the US Court of Appeals in August 2021 regarding the FCC decision to grant license

amendments for SpaceX Starlink satellites [148].

In the SpaceX Intervenor Brief, published in December 2021, the FCC rejected the claim that NEPA requires an environmental assessment of the SpaceX Starlink constellation [149]. Their argument included five main points:

1. Viasat's claims that reentering SpaceX satellites will pollute the atmosphere were "insufficient" and "too vague to warrant further consideration" since the FAA had already concluded SpaceX's launches had "no significant impact" during their own environmental assessment.
2. Risk of human casualty from debris and reentering satellites are "roughly zero" based on the record.
3. Given the "robust record" on reflective satellites and SpaceX's existing mitigation efforts, the FCC determined that an environmental assessment is not necessary, despite Viasat's and The Balance Group's opinion that reflected sunlight on satellites causes "aesthetic, scientific, cultural, social, and health" issues.
4. The FCC has already reviewed SpaceX's orbital debris mitigation plan and finds that it is consistent with existing rules and the public interest. While Viasat claimed that increased collision risk will hamper human development and exploration in space, as well as causing economic harm, the FCC did not think the reasons provided enough detailed justification for further environmental assessment, nor were they convinced that space collisions were even within scope of NEPA, which applies to human environments.
5. Determining that SpaceX already confirmed compliance, the FCC rejected the Balance Group's assertion regarding radiofrequency exposure.

The court ultimately determined that NEPA does not apply to the space environment. This decision highlights the regulatory gaps that leave atmospheric emissions

from space activities largely unregulated. While the FCC determined that an environmental assessment is unnecessary, other US agencies, such as the Government Accountability Office, have acknowledged the environmental risks posed by reentry and atmospheric degradation.

Government Assessment of Environmental Risks

In September 2022, the U.S. Government Accountability Office conducted a technological assessment of large satellite constellations, evaluating the potential environmental side effects of large satellite constellations [150]. They considered emerging technology and strategies to mitigate adverse side effects along with their challenges as well as policies that could address these challenges. Along with orbital debris, they included emissions in the upper atmosphere as an adverse side effect. Noting the uncertain impact of atmospheric emissions, they advocated for more attention and research on the subject in order to guide potential standards, regulations, and agreements. They deemed that not enough is known about emissions in the upper atmosphere to design policies for it.

Observation 20: Existing Policy Proposals

Several proposals exist for policy and regulation to address orbital congestion, but not enough is known about environmental impact to regulate atmospheric emissions. Earth's upper atmosphere is not within NEPA jurisdiction, nor included in the Montreal Protocol.

2.10.3 Policy Mechanisms to Mitigate Orbital Congestion

While there are no existing policy proposals to specifically incentivize OOS in LEO and not enough is known about the impact of atmospheric emissions to regulate it, it is possible that orbital congestion policies intended to incentivize ADR and motivate proper disposal could also create a better business case for OOS. Curbing orbital

Figure 12: Policy framework showing the interrelationship of challenges facing effect mitigations and the context of more persistent uncertainties




Figure 2.15: GAO Policy Framework [150]

congestion and reducing atmospheric emissions share many of the same goals and required capabilities, so policies intended to curb orbital congestion could be adapted for atmospheric pollution reduction.

Policymakers and researchers, intent on establishing policies that will “help, not hurt” their space industries, are exploring different methods to rectify the growing tragedy of the commons without hampering the industry’s growth. Policies intended to incentivize the responsible use of space include orbital use fees, collision charges, mandatory insurance, rebates, deposit/refunds, and launch fees.

Motivating Question 6

What are existing orbital congestion policies, either proposed or implemented, that could incentivize OOS in LEO?

Pigouvian Taxes and Orbital Use Fees

Adilov et al. (2015) propose a two-part Pigouvian tax to entice investment in debris mitigation, finding that without it, investment in debris mitigation falls short of the social optimum while the number of launches far exceeds it [151]. The proposal includes both a launch tax and a tax dependent on the spacecraft operator's preferred debris mitigation strategy. They suggest that revenue from the tax could pay for debris removal but recognize that this tax would be difficult to enact and enforce. Macauley et al. also make the argument for launch taxes that would amount of 0.2 to 2% of production costs [152].

Orbital use fees (OUF) are another approach to combat orbital congestion. Rao et al. recommend an optimal orbital use fee of \$14,900 per satellite that grows to \$235,000 per satellite by 2040 [153]. They argue that not only would OUFs significantly increase the value of the space industry, but they are also necessary to truly mitigate collisions, since satellite operators do not thoroughly account for the risk of collision costs they impose on one another.

Roy et al. consider the role of OUF in promoting on-orbit recycling activities [154]. They assume that the international community would be able to collaborate on establishing OUFs and focus on the impact that the fee would have on operator decision making. Speculating that OUFs would encourage operators to properly dispose of old satellites and a portion of the OUF revenues would finance on-orbit recycling, Roy et al. expect OUFs to generate demand for ADR and recycling. As expected, they found that high OUF shortened break-even time. Overall, the results indicated that profits were more sensitive to space-activity level than OUF policies [154].

Launch taxes, orbital use fees, and debris disposal fees could all contribute to a “polluter pays” superfund that is used to subsidize sustainable operations, such as debris cleanup. In Section 2.10.4, this thesis explores this concept in more detail and

highlights examples from other industries.

Deposit/Refund Systems

Another option for enticing responsible use of space resembles Aldi’s shopping cart strategy. At Aldi’s, you pay a quarter to use the cart and get a quarter back when you return it. The end result: Aldi’s doesn’t have to corral their carts, and you keep your quarter. Macauley (2015) proposes a deposit/refund method with an additional rebate for spacecraft operators who provide a net benefit to the space environment [152]. Their model incorporates collision probability to provide estimates for the deposit, refund, and rebate values.

Insurance-Based Mechanisms

Other policy options include mandatory property insurance and absolute third-party liability insurance in orbit [85]. Currently, some countries only require property and liability insurance during launch, while on-orbit insurance is voluntary. The Convention on International Liability for Damage Caused by Space Objects, established in 1972 and still in effect today, requires proving fault in orbit for insurance claims, which is difficult to do. Consequently, only 5% of current LEO commercial satellites are insured [85].

Making on-orbit insurance mandatory and liability absolute means operators no longer have to jump through this hoop to make an insurance claim in the event of failure. It also creates a business case for on-orbit servicing, where insurers could outsource servicing contracts and reap salvage benefits, such as useful parts and materials.

On-orbit insurance provides a promising option to both mitigate congestion and encourage OOS. In the Aerospace Corporation report “Assurance Through Insurance And On-orbit Servicing,” Reeseman explains that OOS and on-orbit insurance could

have a “unique, symbiotic” relationship [155]. OOS could significantly reduce annual insurance premiums because there would be fewer claims if customers are able to service failed satellites. For instance, Viasat-3 may file an insurance claim worth \$420 million due to its failed reflector deployment, placing a heavy burden on the insurer [156]. Additionally, insurance companies could promote standardization by up-charging satellites that are not service-friendly.

The UK Space Agency has already considered insurance as a tool to promote sustainability. The UKSA consultation report “Orbital Liabilities, Insurance, Charging and Space Sustainability” (2023) recommended removing or reducing the mandatory insurance minimum (EUR 60 million) for operators who conduct sustainable missions and implement practices such as collision avoidance, enhanced trackability, and reduced dark/quiet sky interference [157].

Alternative Policy Approaches

Another proposal to control orbital congestion of the cap-and-trade concept, where governments set a maximum number of allowable space objects, determined based on their space sustainability rating (SSR) and other metrics [158] [159]. Much like a carbon tax, spacecraft operators could trade orbital allowances and earn credits if they deorbit satellites or pay for active debris removal. While this method avoids the risk of undertaxing for very harmful operations in space, it comes with several challenges. It would be difficult to ensure international compliance and it could enable inaction in orbital debris cleanup among the largest satellite operators.

Penalties, like the FCC fine for Dish’s failed disposal, are simple and already demonstrated. Proceeds from these fines could subsidize OOS infrastructure. One downside of this approach is that if 100% of spacecraft operators follow current contractual agreements for proper disposal via deorbiting, there will be no income for OOS. One possible adaptation is to impose a fine when spacecraft operators deorbit

satellites with short life-cycles after a single use. This could further subsidize OOS while encouraging spacecraft operators to choose service.

Several researchers note that governments could subsidize OOS infrastructure while customers pay for the marginal cost of servicing [45] [141]. This approach has been demonstrated before with infrastructure projects such as the Federal Highway Act of 1956, when the government paid for the national highway system. However, the government subsidy approach burdens the taxpayer and is highly sensitive to the political climate and current administration.

Overall, the previous subsections highlight the opportunities for market-based policy to provide course-correction when the private sector is disinclined to improve the sustainability of their operations, which leads us to the following observation. While many of these policy proposals would likely curb congestion in space, they have drawn criticism from companies like SpaceX.

Observation 22: Market-Based Policy Potential

Market-based policies intended to address orbital congestion, such as orbital use fees, taxes, rebates, fines, and required insurance, can internalize environmental costs and motivate private entities to make environmentally-minded decisions.

Ex Ante vs Ex Post Policy Approaches

There are generally two schools of thought when it comes to policies for space debris management. The first is to adapt institutional frameworks that already exist; the other is to apply policy only as a means of promoting the Coase theorem. If there is clarity over property rights, coordinated efforts among operators, and joint research and development, spacecraft operators could negotiate symbiotic terms, despite the externalities.

Another important distinction in policymaking is ex ante versus ex post policies. Ex ante policies make forecasts and anticipate externalities in order to set policy

values, such as taxes and rebates. Meanwhile, ex post policies respond to events as they happen. Ex ante policies are challenging to implement due to the complexity and necessary assumptions involved in developing their forecasting methods.

In September 2023, SpaceX filed a letter with the FCC regarding their position on an aggregate collision probability metric [160]. Promoting the Coase school of thought, they point out that information sharing, coordination, good performance, transparency, and real-time warnings are the most effective means of addressing collision risks and promoting a sustainable LEO environment. Addressing Astroscale and OneWeb’s proposed aggregate collision probability (ACP) metric, an ex ante measure, SpaceX makes six key points:

1. There is no universally accepted aggregate collision probability calculation.
2. The ACP metric compounds conservative estimates that ultimately punish spacecraft operators in LEO.
3. The ACP metric doesn’t address the “true” risks to space sustainability because it de-emphasizes the GEO regime, where space debris takes far longer to deorbit.
4. The ACP method would harm U.S.-licensed systems, prompting forum shopping and benefiting foreign operators.
5. The ACP metric would harm American consumers because it would disproportionately disenfranchise the LEO economy, which is inherently safer than GEO due to its “self-cleaning” nature.
6. The proposed “safe harbor” could potentially make the problem worse, because while the responsible, incumbent operators are penalized, the new, less responsible operators get free rein.

SpaceX condemns policy based on ex ante analysis and overly complex methodology. Amidst the tense socio-technical space regime, universal policy buy-in is more

likely when the metrics are transparent and follow the KISS principle (Keep It Simple, Space-policymaker).

The various policy mechanisms explored in this section share a common economic foundation: they aim to internalize the environmental externalities that satellite operators currently impose on the space environment without bearing the associated costs. When private actors make decisions based solely on their individual costs and benefits and ignore the broader environmental consequences of their actions, a tragedy of commons emerges. This coordination failure means that individually rational decisions lead to collectively suboptimal outcomes for the space environment and the long-term viability of space activities.

Policy intervention becomes necessary when market forces alone fail to account for these environmental costs. However, effective policy design requires careful calibration to avoid unintended consequences. Well-designed policies should expand economic opportunities while correcting market failures, rather than simply imposing punitive measures based on imprecise metrics or arbitrary thresholds. The goal is to align private incentives with environmental sustainability, creating win-win scenarios where profitable business decisions also support long-term space environmental health.

Motivating Question 7

How to design policy to control orbital congestion and mitigate atmospheric degradation without causing unnecessary harm to the industry?

2.10.4 Policy Mechanisms from Analogous Industries

Space is hardly the first industry to suffer from a tragedy of commons. To inform the space policy conversation and gain a better understanding of what policies could alter behavior without penalizing the industry, it is worthwhile to explore the policies and regulations applied to other fields and consider their effectiveness.

Polluter-Pays Principle

The polluter-pays superfund concept has been gaining increasing attention among space policy researchers, with many suggesting that it could hold the key for incentivizing sustainable operations in orbit [161] [162] [159]. In “Confronting Space Debris”, Baiocchi and Welser of the RAND Corporation consider 9 issues, ranging from hazardous waste to email spam, that share characteristics of the orbital debris problem [161]. They considered three mitigation concepts, command and control, market-based approaches, and performance-based approaches such as quota-systems. Based on their compiled research, they concluded that a superfund is a promising method to remediate orbital debris. They also found that incentivization methods are effective in the short term, but that in order to arrive at a long-term solution, stakeholders must change their priorities.

The polluter-pays principle (PPP) is an ancient one, with roots dating back to Plato and his recommendations for water pollution laws within his *Nomoi* (Laws) [163]. Around the same time in the 3rd century BCE, the ancient Indian philosopher Kautiliya defines various financial penalty levels for those who harm the environment, dependent on their degree of harm, within the *Arthashastra* (Study of Economics) [164]. The polluter-pays principle now inspires or forms the basis for a wide range of international environmental agreements, referenced in the Paris Climate Agreement and the Helsinki Principle [165]. The major US pollution laws, such as the Clean Air Act, the Clean Water Act, the Resource Conservation and Recovery Act, and the Superfund borrow elements of the polluter-pays principle, but do not apply the principle to the full extent. The Superfund, for instance, which is a fund to address hazardous waste cleanup, charges polluters based on their ability to pay rather than the true extent of their damages.

There are a number of other policies that apply the polluter pays principle to some extent. In 2024, both Vermont and New York State established climate superfund

programs that require fossil fuel companies to contribute to climate adaptation funds. Vermont's Climate Superfund Act, which became law in May 2024, allows the state to recover financial damages from fossil fuel companies for climate change impacts [166] [167]. Loosely based on the federal Superfund program (CERCLA), the law requires companies responsible for more than one billion tons of greenhouse gas emissions globally to make payments based on their emissions between 1995 and 2024 [168] [169]. New York followed suit when Governor Kathy Hochul signed similar legislation in December 2024, requiring large fossil fuel companies to pay for critical climate protection projects [170] [171]. While other states, such as Maryland, Massachusetts, and California are working on their own version of the bill, both New York and Vermont state laws are likely to face legal challenges based on federal preemption and due process [169]. Consequently, it may be difficult to measure the success of these programs, but they signal growing interest in superfunds to address environmental issues.

The Corporate Average Fuel Economy (CAFE) standards operate as a polluter pays mechanism by regulating vehicle fuel efficiency, with manufacturers historically required to pay penalties when their vehicles fall short of the standards. These penalties increased from \$5.50 to \$14 for every 0.1 mpg that new vehicles failed to meet standards for model years 2019-2021, with the fine calculated by multiplying this rate by the number of non-compliant vehicles sold [172]. However, in July 2025, Congress eliminated civil penalties for noncompliance with federal fuel economy standards for passenger cars and light trucks, though the CAFE statute itself remains on the books and companies still have reporting obligations [173]. This elimination of penalties has changed manufacturers' incentive structures and reduced the value of CAFE compliance credits that electric vehicle manufacturers had previously relied upon as a source of revenue by selling them to less fuel-efficient manufacturers [173].

The European Union Emissions Trading System applies the polluter pays principle

by requiring companies to purchase emissions allowances for each tonne of carbon dioxide they release into the atmosphere [174]. The system operates on a cap-and-trade principle where companies must monitor and report their annual emissions and surrender sufficient allowances to cover them, with the ability to trade allowances among themselves as needed [175]. Recent reforms have accelerated implementation of the polluter pays principle by phasing out free allowances for the aviation sector, moving toward full auctioning by 2026 [176]. This market-driven mechanism creates funding for cleaner alternatives by putting a price on pollution.

There are several challenges in creating a PPP policy, such as ensuring that the cost polluters pay is commensurate with their damage, which is often difficult to measure. There's also the question of how to hold accountable entities whose past pollution still causes harm today, even though they are actively making efforts to reform, compared to entities actively worsening the problem, which is illustrated by current climate change debates. While the PPP concept is simple — if you damage it, you pay for it — execution at the international scale becomes prohibitively complex. A more tractable approach is small, national-level subsidy fund where governments tax their domestic satellite companies to create modest funds that provide just enough subsidy to make sustainable space operations economically viable. By closing the business case gap, these smaller-scale programs could drive change with minimal protests from the private sector and without requiring international coordination. If successful, these national-scale superfunds could lead by example and prompt other governments to follow suit. When it comes to building and utilizing a superfund to subsidize or reward sustainable operations, there is more than one way to peel a carrot. The following sections describe these various types of implementations.

Observation 23: The National-Level Subsidy Fund

While not without their challenges, superfund programs, particularly those at the regional/national level to avoid issues with international compliance, have proven effective at promoting private sector environmental stewardship through economic incentives rather than prescriptive regulations.

Feebate System

A feebate system represents a self-contained policy mechanism where a government taxes an industry to create a fund that subsequently finances incentives for environmentally beneficial actions within that same industry. This circular funding structure internalizes the costs of pollution while simultaneously subsidizing cleaner alternatives, creating a revenue-neutral approach that encourages industry transformation without requiring broader taxpayer funding.

The fundamental principle of a feebate system is straightforward: entities engaging in environmentally harmful practices pay fees that directly fund rebates or subsidies for entities within the same industry that adopt more sustainable practices. This creates a competitive advantage for cleaner operations while penalizing polluters, effectively using market mechanisms to drive environmental improvements. The system is considered revenue-neutral because the fees collected match the rebates distributed, maintaining fiscal balance while reshaping industry incentives.

Several real-world implementations demonstrate the versatility of feebate systems across different sectors, both on large and small scales:

- **Renewable energy certificates trading:** Fossil fuel companies purchase credits from renewable energy producers, creating a direct financial flow from polluting energy sources to clean alternatives within the energy sector [177].
- **Plastic bag levy systems:** Fees imposed on single-use plastic bags fund programs promoting reusable bags and broader waste reduction initiatives, trans-

forming consumer behavior through financial incentives derived from the problematic practice itself [178].

- **Carbon pricing with revenue recycling:** Several jurisdictions implement carbon taxes or cap-and-trade systems with explicit revenue recycling provisions. British Columbia's carbon tax, introduced in 2008, returned all revenue to residents and businesses through tax cuts, creating a revenue-neutral system that changed behavior while reducing burden on low-income households [179]. Within 5 years of implementation, the tax reduced fuel use by 16% and emissions are estimated to have reduced by 15% [179]. Economic studies found no negative financial impact and even some small, positive impact as well as job growth.

These examples illustrate how feebate systems create closed-loop economic incentives that drive environmental improvements without requiring external funding sources, making them particularly attractive for politically sensitive contexts where new taxes or government expenditures face resistance.

Sectoral Transformation Fund

A sectoral transformation fund represents a more comprehensive policy mechanism than a simple feebate system, combining both behavioral incentives and infrastructure development. Under this approach, a government taxes an industry to create a fund with a dual purpose: first, to provide direct payments or rebates to entities within that industry when they adopt environmentally beneficial practices; and second, to subsidize the development of sustainable infrastructure that enables broader industry transformation. This dual-pronged strategy addresses both immediate behavioral change and long-term structural transformation, creating an internalized mechanism for industry modernization.

Fuel taxes generate substantial revenue for the Highway Trust Fund (HTF), which supports highway and transit infrastructure. The HTF, established in 1956, represents one of the oldest examples of dedicated sectoral funding [180]. While these funds primarily support infrastructure development, such as roads, bridges, and public transit systems, rather than direct incentives for fuel-efficient vehicles, they demonstrate the political durability and fiscal stability of dedicated revenue streams that remain within the transportation sector.

Sectoral transformation funds combine infrastructure subsidies with some of the mechanisms of feebate systems. Infrastructure investments create network effects and economies of scale that amplify the impact of individual behavioral incentives. For instance, electric vehicle adoption incentives become more effective when paired with public investment in charging infrastructure. Additionally, subsidized sustainable infrastructure addresses the Tragedy of Commons issue and provides a demand signal to the private sector that the government is committed to transformations in sector sustainability, providing certainty that encourages private sector planning and investment in complementary sustainable technologies.

Sectoral transformation funds operate through a circular funding mechanism where the industry essentially finances its own modernization. Tax revenue collected from environmentally harmful activities within the sector creates a dedicated fund that cannot be diverted to general government purposes. This earmarked structure ensures that polluters directly subsidize both the transition of their competitors to cleaner practices and the infrastructure investments that benefit the entire sector. The revenue neutrality of this approach, where funds stay within the industry rather than flowing to general taxation, makes it politically more palatable than broader environmental taxes while still achieving substantial environmental outcomes.

There are several examples that demonstrate the effectiveness of sectoral transformation funds across different industries:

- **Regional Greenhouse Gas Initiative (RGGI):** RGGI, with 12 states participating in the US, is one of the most comprehensive examples of a sectoral transformation fund in practice [181]. RGGI states auction carbon dioxide emission allowances from fossil fuel power plants, generating substantial revenue that is reinvested within the energy sector. Between 2009 and 2017, RGGI produced a net benefit of \$4.7 billion through the cap-and-trade program [181]. These funds are deployed through a dual strategy: direct incentives for energy efficiency and clean energy adoption by households and businesses, alongside major infrastructure investments in grid modernization, renewable energy systems, and charging infrastructure [182]. Saving 7.8 million tons of CO₂, the lifetime benefits of investments made in 2023 alone (\$2.7 billion) demonstrate the model's effectiveness at driving both individual behavioral change and systemic transformation [182].
- **Waste disposal surcharges:** Many state and local jurisdictions levy per-ton fees on waste sent to landfills or incinerators, creating dedicated funds that support comprehensive waste diversion programs [183]. These surcharges fund both recycling incentives that reward individuals and businesses for diverting waste, and the construction and operation of recycling and composting facilities that enable those sustainable practices. Indiana's programs awarded over \$1.8 million in grants that increased recycled materials by 85,000 tons and created 47 jobs, while Pennsylvania's Food Recovery Infrastructure Grant provided \$9.6 million to projects that rescued nearly one million pounds of food and served over 25,000 residents [183]. These programs demonstrate how disposal fees can simultaneously change behavior and build the infrastructure necessary for system-wide transformation.
- **Electronic waste recycling fees:** States like California (The Electronic Waste

Recycling Act) impose fees on purchases of electronic devices, creating funds dedicated to electronic waste recycling infrastructure and consumer take-back programs [184]. These fees support both the convenient recycling opportunities that encourage proper disposal and the specialized facilities required to safely process electronic waste [185].

These examples illustrate varying degrees of integration between behavioral incentives and infrastructure development. RGGI and waste disposal surcharges most closely embody the sectoral transformation fund concept, combining direct incentives for sustainable practices with infrastructure investments that enable broader adoption. Together, these cases show that the combined approach of behavioral incentives and infrastructure subsidies create a more powerful transformation dynamic than either component alone, accelerating industry change through mutually reinforcing mechanisms.

The combined approach of behavioral incentives and infrastructure subsidies creates a more complete cycle of transformation. Initial infrastructure investments reduce the cost and increase the convenience of sustainable practices, making behavioral incentives more effective in driving adoption. As adoption increases, economies of scale further reduce costs, making additional infrastructure investments more economically justified. This feedback loop accelerates the pace of industry transformation beyond what either component could achieve independently, making sectoral transformation funds particularly well-suited for achieving ambitious environmental goals within specific industries while maintaining political feasibility through revenue neutrality and direct industry benefits.

2.10.5 Policy Design for Closing the Business Case Gap

The fundamental challenge for OOS policy design lies in addressing the business case gap identified earlier: while OOS may generate positive net benefits at the system

level, individual operators face insufficient private incentives to invest in servicing infrastructure or design serviceable satellites. Traditional regulatory approaches, such as comprehensive environmental frameworks with strict mandates, risk stifling innovation and imposing significant compliance costs on an emerging industry.

A more targeted approach focuses on surgical policy interventions that specifically address the coordination failure between satellite-operator and service-provider underlying the OOS business case gap. Rather than broad regulations, carefully calibrated market-based mechanisms can shift competitive dynamics by making OOS economically attractive relative to traditional disposal and replacement strategies. The key insight is that relatively modest policy interventions can tip the balance when the underlying economics are already close to viability.

Previously, this thesis identifies collection hubs that provide collection-as-a-service and flexible options as pathways toward de-risking and implementing on-orbit servicing in LEO. Policy can play a crucial role in this goal as well, providing a catalyst that de-risks infrastructure deployment to improve returns over the long-term horizon compared to laissez-faire configurations where satellite constellation operators make deployment decisions with neither stick nor carrot. By designing policy interventions that interplay with flexible deployment strategies, it becomes possible to motivate circularity while minimizing the regulatory burden on constellation operators.

The most promising policy designs share a common characteristic: they can be structured to be approximately revenue-neutral from the satellite-operators' long-term perspective while still creating powerful incentives for sustainable behavior. This approach addresses political feasibility concerns while ensuring that policy interventions genuinely solve coordination problems rather than simply redistributing costs within the industry.

Observation 24: Revenue-Neutral Policy and the Sectoral Transformation Fund

Rather than comprehensive environmental regulation, targeted policy mechanisms can “close the gap” for the OOS business case. Policy interventions, such as the sectoral transformation fund concept, can be designed to be approximately revenue-neutral while still shifting competitive dynamics to favor OOS, minimizing political resistance and industry burden while addressing the coordination failure.

Policy Scheme Overview

While numerous policies could influence the OOS market in LEO, this research examines eight distinct policy approaches that represent different mechanisms for internalizing environmental costs and addressing the coordination failures identified in the business case gap. Each scheme targets the common goal of making on-orbit servicing economically attractive while minimizing regulatory burden and industry disruption.

The eight policy schemes tested in the flexibility framework are:

- **Policy Scheme 1:** Orbital Use Fee with Refund implements a deposit/refund system where operators pay annual fees for each satellite but receive full refunds when satellites are either collected and brought to servicing infrastructure or refurbished on-orbit. This approach directly incentivizes responsible end-of-life management and satellite servicing.
- **Policy Scheme 2:** Orbital Use Fee with Subsidy collects annual fees and uses the proceeds to subsidize warehouse purchases, infrastructure upgrades, and refurbishment rebates. This addresses the collective action problem where individual operators lack incentive to invest in environmentally-oriented infrastructure.
- **Policy Scheme 3:** Contingent Fines with Subsidy operationalizes a “polluter

pays” framework by fining failed satellites while providing refunds for operators who collect failed assets. Proceeds subsidize servicing infrastructure and provide refurbishment rebates, creating direct accountability for satellite failures while maintaining economic incentives for remediation activities.

- **Policy Scheme 4:** Mandatory Insurance requires operators to carry on-orbit insurance for each satellite, but waives premium requirements for satellites that receive servicing. This leverages existing insurance market mechanisms while creating direct economic incentives for utilizing available servicing infrastructure.
- **Policy Scheme 5:** Progressive Taxation with Subsidy implements time-progressive taxes on constellation operator profits, with proceeds subsidizing sustainable infrastructure. This approach scales obligations with operator success rather than imposing uniform burdens regardless of financial circumstances.
- **Policy Scheme 6:** Subsidy/Taxes/Fine combines the progressive taxation approach (Policy 5) with the contingent fine mechanism (Policy 3), creating a dual-incentive structure that both penalizes failures and scales contributions to operator profitability while funding infrastructure and refurbishment rebates.
- **Policy Scheme 7:** Subsidy/OUF/Fine integrates orbital use fees (Policy 2) with contingent fines (Policy 3), combining baseline fees on all satellites with additional penalties for failures. Proceeds fund infrastructure subsidies and refurbishment rebates, creating layered economic incentives for sustainable operations.
- **Policy Scheme 8:** Subsidy/OUF with Premium implements a premium-based orbital use fee structure (similar to mandatory insurance) rather than a flat annual fee, with proceeds subsidizing infrastructure purchases, upgrades, and

refurbishment rebates.

The detailed implementation, parameter selection, and modeling methodology for these schemes are described in Chapter 3. The flexibility framework enables systematic comparison of these policy approaches against laissez-faire baselines to evaluate their effectiveness in closing the OOS business case gap.

2.10.6 Research Gaps and Framework Development

The policy landscape analysis reveals three critical research gaps that prevent effective policy design for incentivizing OOS in LEO. These gaps collectively demonstrate the need for a systematic framework to evaluate policy interventions and improve their parameters for maximum effectiveness with minimal industry disruption.

Research Gap 6: Environmental Policy Gaps

No policy proposals exist specifically to limit atmospheric pollution from space activities or to incentivize OOS in LEO through circular economy approaches.

While orbital congestion policies have received significant attention in the literature, the specific challenge of incentivizing circular economy approaches in space through OOS remains unaddressed. Existing congestion policies could potentially be adapted to address atmospheric emissions concerns, but no research has systematically explored this connection.

Even if appropriate policies were proposed, policymakers currently lack tools to evaluate how different interventions would affect the economics of servicing infrastructure deployment and operation. This analytical gap means that well-intentioned policies could inadvertently harm rather than help the development of sustainable space infrastructure. The complex interplay between policy incentives, infrastructure investment decisions, and operator behavior requires modeling approaches that account for uncertainty and flexible deployment strategies.

Research Gap 7: Policy-Infrastructure Interaction

No flexibility frameworks exist that measure the impact of policy intervention on infrastructure cost/benefits or investigate how congestion policies could impact demand for OOS in LEO.

The parameter design challenge represents a critical gap for practical policy implementation. Even with appropriate policy mechanisms identified, determining optimal fee levels, rebate structures, and implementation timelines requires quantitative analysis that balances multiple objectives: closing the business case gap, minimizing industry burden, maintaining revenue neutrality, and avoiding unintended market distortions. Without systematic parameter exploration, policies risk being either ineffective at incentivizing change or unnecessarily burdensome to industry stakeholders.

Research Gap 8: Policy Parameter Design

No frameworks exist to identify policy parameters (orbital use fee levels, rebate structures, progressive fee designs, etc) that close the OOS business case gap in LEO while reducing market and regulatory burden.

These three research gaps collectively point towards a fundamental need for integrated policy design and evaluation frameworks. This need gives rise to the central research question for policy intervention:

Research Question 4: Policy Design

Which combination and calibration of policy parameters (annual orbital use fees, collection/servicing rebates, progressive fee structures) most effectively closes the business case gap for OOS in LEO with minimal impact on overall industry costs?

This research question addresses all three identified gaps by requiring: (1) development of specific policy proposals for OOS incentivization, (2) creation of frameworks

to measure policy impact on infrastructure economics, and (3) systematic design space exploration of policy parameters. Answering this question requires a modeling framework capable of comparing different policy approaches under uncertainty while accounting for the flexibility inherent in infrastructure deployment decisions. Adapting the flexibility framework to accommodate the costs and benefits associated with various policy mechanisms and their parameters provides a test bed for addressing this question, leading to the third and final hypothesis.

Hypothesis 3

If the flexibility framework models government intervention as parametric, time-dependent “rewards” and “penalties”, then there exist scenario-dependent reward/penalty schemes that best establish economically feasible OOS infrastructures that yield better sustainability metrics than laissez-faire OOS infrastructures developed with the same incremental deployment framework

This hypothesis enables systematic testing of the eight policy schemes introduced earlier by integrating them into the flexibility framework as parametric interventions that can be varied and compared. The framework provides a testbed to compare infrastructures developed laissez-faire, without policy, with those developed under various policy interventions. By modeling policies as time-dependent parameters that affect the costs and benefits of different operator decisions, the framework can identify top-performing policy combinations and calibrations for different scenarios and system objectives.

The testing of this hypothesis directly addresses the identified research gaps by creating the first systematic flexibility framework for evaluating OOS-specific policies, measuring their impact on infrastructure development, and tailoring their parameters for maximum effectiveness. This approach moves beyond theoretical policy proposals toward practical policy design tools that can inform real-world decision-making.

2.11 Thought-Experiment and Real-World Analogies

Collecting old technology that is abandoned in space? Making a business of rehabilitating and refurbishing salvage? Science fiction aficionados will recognize that this is not a new concept, having been perfected by the Tatooine dwellers, Jawas. Letting nothing go to waste, Jawas create something of a circular economy by collecting old droids, refurbishing them, and selling them for profit. Tatooine’s desert would be littered with droid parts if not for these opportunistic environmentalists.

While Jawas provides an entertaining analogy, real-world examples of circular economy principles offer concrete insights for space operations. The transition from linear “take-make-dispose” models to circular systems has been successfully demonstrated across multiple terrestrial industries, each providing valuable lessons for satellite lifecycle management.

Distributed Service Networks Real-life analogs for satellite collection exist in distributed service networks, such as the Gogoro electric scooter battery swap stations in Taiwan [186]. The Gogoro scooters are quickly gaining popularity in East Asia because people can easily swap their dead battery for a charged one at the charging stations, which are located at convenient locations. The positioning of these charging stations is analogous to the positioning of the collection hub/spare warehouses in a parking orbit that drifts through orbital planes. It would be wasteful to throw the scooter battery out after it loses its first charge, just as single use satellites fail to fulfill their full value. Charging the scooter battery at home and waiting for it to charge before you can use it again is inconvenient, just like waiting for a mobile servicer to reach a satellite at the end of its lifecycle or after a random failure. It is much more convenient to replace the satellite with a readily available and refurbished spare.



Figure 2.16: Real World Analogies of Circular Systems: Gogoro Battery Station (left) [186] and Daisy Robot (right) [187]

Automated Disassembly and Refurbishment Analogies also exist to support the idea of refurbishing satellites on Earth. Apple’s disassembly robot, called Daisy, is capable of taking apart 29 different types of iPhones. Processing each iPhone within 18 seconds, it can process around 1.2 million iPhones each year [187]. While this is a relatively small percentage of the 150 million iPhones thrown out in 2023, it marks the first steps towards improved system circularity. In the same vein, if satellites were brought back to Earth for recycling, satellite operators would have a reason to develop streamlined refurbishing processes and revisit satellite design to make them more easily refurbishable, even if it’s a relatively small percentage of the overall constellation. These developments would then be transferable to orbit in the future, making a down payment towards improved circularity in space.

Broader Circular Economy Precedents Similar transformation patterns emerge across other industries. The automotive sector has developed sophisticated re-manufacturing processes for engines, transmissions, and electronic components, often achieving performance equivalent to new parts at 30% of the original cost [188]. The telecommunications industry is also working on refurbishing and redeploying cellular infrastructure equipment, extending operational lifespans through component upgrades and repairs. Engineers at Ericsson have determined that refurbishment reduces supply-chain re-

lated carbon emissions by >90% compared to manufacturing a new product [189]. In the US, researchers and policymakers are working on better ways to streamline and incentivize battery recycling [190]. While there are several challenges in recycling and reusing batteries, recovered and reintegrated battery materials would reduce both resource extraction and waste streams.

These terrestrial examples share common characteristics that make them relevant to space operations: high-value assets with modular designs, established reverse logistics networks, standardized refurbishment processes, and economic incentives that make circular approaches competitive with linear alternatives. Each industry's evolution toward circularity was driven by a combination of resource scarcity, regulatory pressure, and economic opportunity—factors increasingly present in the space domain.

2.12 Overarching Research Framework

Overarching Observation

The process of motivating and establishing circular space economies in LEO is a complex system of systems problem that requires analysis from the technical, financial, and policy perspectives. This thesis aims to provide a systems-level screening framework that evaluates the interaction between novel OOS CONOPs, flexible options, and various policy schemes in order to path-find strategies and infrastructures that could improve the case for OOS in LEO.

The overall thesis objective is to identify opportunities that sufficiently incentivize the private sector to improve sustainability (via circularity) in LEO and reduce atmospheric emissions associated with the operations and maintenance of constellations composed of single-use satellites. This chapter has provided an exploration of CONOPs, policy, and business strategies that could improve the case for OSAM in LEO, taking into account the numerous sources of uncertainty. Flexibility frame-

works that include novel and state-of-the-art OSAM CONOPs will provide a means of investigating options that improve the LEO OSAM business case. The overarching research question is as follows:

Overarching Research Question

Which flexible option, or set of options, consistently improves the economic value and environmental impact of LEO-based OOS over a range of potential policies and future scenarios sampled from multi-domain uncertainty?

By applying the impact of policy parameters to the cost and benefits of mega-constellation operators and OOS providers, the framework provides a computational laboratory for investigating different policies and their impact on demand for OOS, which will provide a better understanding of how policy could influence the state of the circular space economy. This thesis posits that novel OOS concepts, combined with flexible strategies and policies, will enable economically feasible OOS infrastructures.

Overarching Hypothesis

If a flexibility framework for LEO-based OOS incorporates multiple uncertain variables, policy impact, novel design concepts like collection hubs, and allows for multiple combinatorial options, then there will exist an option or set of options that provides a viable and sustainable private infrastructure for circular space economy.

To substantiate this overarching hypothesis, there needs to be an assessment methodology that allows the user to screen the value of flexible options over a wide variety of scenarios. This leads us to the Overall Framework Gap:

Overall Framework Gap

No existing flexibility framework for LEO-based OOS includes multi-domain uncertainty, multi-agent decision making, and the combinatorial effects of both classic OOS options and novel concepts (cooperative maneuvering, collection hubs, temporary abandonment) that would allow screening for economically and environmentally feasible strategies and policies.

This chapter contains an evaluation of the research gaps and presents the hypothesis. The following chapter presents the specific methodology and framework details in order to substantiate the presented hypotheses.




Figure 2.17: Framework Gap for Experimentation Conditions

CHAPTER 3

METHODOLOGY

The methodology developed in this chapter draws from analogous applications and infrastructure projects while addressing the unique constraints of space operations in LEO. The following sections provide details on the specific CAAS CONOPs, identify necessary simulation properties, identify and characterize sources of uncertainty, develop the flexibility framework, model the effect of various policy options, and provide a methodology for path-finding options and strategies that could enable the space industry to improve their space junk allocation and waste abatement (Space JAWA).

3.1 CONOPs Summary

The Collection-as-a-Service (CAAS) CONOPs represents an advanced multi-echelon sparing strategy designed for mega-constellations in Low Earth Orbit (LEO). This framework extends traditional satellite constellation management by incorporating Active Debris Removal (ADR) vehicles, orbital collection hubs, and cooperative maneuvering to create a flexible, sustainable infrastructure for satellite operations. The CAAS CONOPs builds upon previously proposed multi-echelon sparing, which features warehouses of spare satellites located in parking orbits that incrementally resupply in-plane spare satellites that are capable of replacing failed satellites in a timely manner. In some cases, multi-echelon sparing may offer cost benefits compared to other sparing strategies, like overpopulation or direct resupply. Figure 3.1 below provides a comparison of these 4 CONOPs. The Collection-As-A-Service CONOPs is summarized as follows:

CAAS CONOPS Overview

System Description: The Collection-as-a-Service (CAAS) framework is an advanced multi-echelon sparing strategy that integrates orbital spare warehouses and Active Debris Removal (ADR) vehicles to provide comprehensive satellite collection, processing, and replacement services for mega-constellations in Low Earth Orbit (LEO).

Key Operational Constraints & Assumptions:

- **Orbital Configuration:** All warehouses and constellation satellites operate at the same inclination. Warehouses are positioned in lower parking orbits where J2 perturbation effects cause natural drift between orbital planes—no propulsive plane changes are used.
- **Economic Model:** The analysis assumes a single entity operates both the constellation and the CAAS infrastructure (warehouses and ADR vehicles), enabling integrated decision-making for optimal system-wide economics.
- **Failure Collection Response:** Warehouse-based ADR deployment depends on urgency of the failure type and several conditions. If satellite failure requires immediate replacement, Warehouse-based ADR deployment occurs only if a warehouse is already in close proximity and can reach the failed satellite within the same timeframe as a ground-launched vehicle. If satellite failure does not require immediate replacement, Warehouse-based ADR deployment occurs if a warehouse is within five years of reaching the failed satellite. Otherwise, an ADR vehicle is launched from Earth.
- **Spare Satellite Inventory:** Two in-plane spare satellites are maintained

per orbital plane across all CONOPS for immediate failure response.

- **Reusable Launch Vehicle Operations:** Reusable second stages return to Earth after delivering satellites to orbital planes. These stages can opportunistically rendezvous with nearby warehouses to collect satellites for Earth return *only if the timing and trajectory permit – no propulsive plane changes*. Rockets are never launched empty (“deadheading”) solely to retrieve satellites.
- **Satellite Maneuvering Capabilities:** If upgraded to do so, operational satellites can maneuver close to warehouses, where ADR vehicles complete final docking and collection. Full autonomous rendezvous and proximity operations with docking (RPOD) capabilities are assumed cost-prohibitive for constellation satellites.

Core Operational Principle: CAAS leverages natural orbital mechanics (J2 drift) and opportunistic operations to minimize propulsive maneuvering costs while maintaining constellation coverage through flexible, multi-pathway satellite collection and processing options.

Baseline Strategy Comparison

Two Baseline Strategies Compared to CAAS:

1. 0-Warehouse Baseline (Overpopulation):

- *Architecture:* Two in-plane spare satellites per orbital plane
- *Spare Replenishment:* Direct launch from Earth to orbital plane when spare is used
- *Failed Satellite Collection:* ADR vehicle launched from Earth for

every failure

- *Key Characteristic:* Simplest architecture with no orbital infrastructure; relies entirely on ground-based launch responsiveness

2. Multi-Echelon Sparing (No CAAS Features):

- *Architecture:* Two in-plane spare satellites per orbital plane + parking orbit warehouses storing additional spares
- *Spare Replenishment:* When in-plane spare replaces failed satellite, warehouse drifts via J2 to replenish in-plane spare inventory
- *Failed Satellite Collection:* ADR vehicle launched from Earth for every failure
- *Key Characteristic:* Reduces launch frequency for spare replenishment through orbital warehouses, but does not integrate ADR with warehouse operations or enable flexible collection timing

Critical Distinction: The multi-echelon baseline provides only the warehouse-based spare replenishment benefit without CAAS operational innovations.

3.1.1 Core Infrastructure Components

The CAAS system operates through a network of orbital spare warehouses positioned in lower parking orbits that gradually drift between orbital planes via J2 perturbation effects. These collection hubs serve multiple functions as both spare satellite storage facilities and processing centers for collected end-of-life spacecraft. The warehouses can be incrementally upgraded with enhanced capabilities includ-




Figure 3.1: Comparison of sparing strategies for satellite resupply.

ing fuel depots, robotic servicing equipment, and spare parts inventory to provide on-orbit refueling and repair services, as shown in Figure 3.2.

Active Debris Removal vehicles form the operational backbone of the collection service, functioning as specialized spacecraft capable of rendezvous, capture, and transport operations. These ADR vehicles operate with significant deployment flexibility as they can be launched directly from Earth for immediate collection missions or deployed from orbital warehouses if a failed satellite is nearby. After completing a mission, they have the option to rendezvous with a passing warehouse to deposit its collection of satellites, resupply on fuel, and prepare for the next collection mission.

Upgrade Option:




Figure 3.2: CAAS Upgrade Option

These options are compared in Figures Figure 3.3 and Figure 3.4.

3.1.2 Operational Framework

The CAAS CONOPs addresses fundamental satellite constellation operator needs by maintaining consistent coverage while reducing operational costs. When satellites experience failures or reach their intended end-of-life, operators have multiple response options within the framework. The system maintains in-plane spare satellites for immediate replacement of failed assets (2 in-plane spares for all CONOPs), while orbital warehouses provide a secondary tier of spare inventory that drift through planes using natural orbital mechanics.

For end-of-life satellite management, the CAAS framework requires constellation operators to utilize ADR services for all failed satellites to reduce collision risk and maintain regulatory compliance for replacement authorizations. However, the timing and urgency of collection operations varies significantly based on operational circumstances. Some satellite failures leave the satellite completely inoperable and unable to

Direct ADR Launch Option:




Figure 3.3: ADR Launched from Earth

maneuver itself, necessitating immediate ADR collection since the satellite may pose a collision threat to its replacement. In contrast, other failures can be immediately replaced using available in-plane spares, allowing the failed satellites to remain in orbit temporarily until ADR vehicles can collect them during scheduled operations. Since the FCC requirement is to dispose of satellites within 5 years of their EOL, the framework allows five years for an ADR vehicle to drift into the vicinity of the failed satellite and collect it. This flexible timing approach ensures continuous service while optimizing collection efficiency.

The CAAS system maximizes operational efficiency through opportunistic rendezvous operations when vehicles pass near orbital warehouses, illustrated in Figure 3.5. If reusable rocket stages that have delivered new satellites to an orbital plane can rendezvous with nearby warehouses during their return trajectory to collect satellites designated for Earth return, they reduce the need for dedicated collection missions. Similarly, ADR vehicles launched directly from Earth to collect specific

Parking Orbit ADR Option:




Figure 3.4: ADR Deployed from Parking Orbit

failed satellites can rendezvous with warehouses along their return path if proximity allows. These ADR vehicles can deposit their collected satellites at the warehouse for processing and subsequently join the orbital fleet to conduct additional collection missions, incrementally expanding the warehouse's ADR capacity without requiring separate deployment launches.

3.1.3 Collection and Processing Options

Once satellites are collected by ADR vehicles, the CAAS system provides multiple pathways for processing based on satellite condition, technological obsolescence, and warehouse capabilities. Collection hubs can accommodate various satellite states and offer different processing options depending on available infrastructure and economic considerations.

Collecting satellites via ADR vehicles is not the only method for wrangling satellites within the framework. Satellites also have the option of being upgraded so that

Plane Resupply with Rendezvous Option:




Figure 3.5: Plane Resupply with Rendezvous Option

they are capable of rendezvous proximity operations to the warehouses, an option that is depicted in Figure 3.5. Once in range, ADR vehicles can collect the satellites and conduct last-mile maneuvering and docking. Upgrading satellites for full RPOD capabilities is assumed to be inhibitively expensive given their price point. RPO, meanwhile, is not a large jump from their existing deorbiting capabilities. Allowing satellites the chance to deliver themselves allows ADR vehicles to focus on collected failed satellites that are unable to maneuver themselves.

Satellites that remain technologically current and require only minor repairs can be directly serviced at warehouses upgraded to provide repair and refuel services. These spacecraft can be refueled, repaired, and immediately recommissioned as operational spare satellites without leaving the orbital environment, maximizing efficiency and minimizing costs.

For satellites requiring more extensive refurbishment beyond the warehouse's on-orbit capabilities, the system can return them to Earth for comprehensive overhaul and technological upgrades, depicted in Figure 3.6. Once refurbished terrestrially,

these satellites can be relaunched and reintegrated into the constellation or spare inventory. All satellite-return missions utilize resupply launches — either delivering new satellites to nearby orbital planes or restocking the warehouse. Rockets are never launched empty solely to retrieve satellites and return them to Earth, a practice known in terrestrial logistics as “deadheading.”




Figure 3.6: Refurbishment on Earth Option

Satellites that have become technologically obsolete or are beyond economical repair have several disposal pathways available. They can be deorbited or returned to Earth for component recovery and material recycling. The choice among these options depends on satellite condition and available infrastructure capabilities.

3.1.4 Flexibility and Adaptability

The CAAS CONOPs incorporates significant flexibility mechanisms to adapt to uncertain future scenarios including varying failure rates, technological obsolescence patterns, and market conditions. The modular nature of the infrastructure allows for incremental expansion or modification based on operational experience and changing

requirements.

The framework addresses the challenge of technology obsolescence by enabling scheduled replacement of orbital spares before they become outdated, treating obsolete spares as candidates for collection and processing. This approach ensures that the spare inventory remains technologically current while providing additional opportunities for satellite reuse and recycling.

Service pricing within the CAAS framework adapts dynamically based on satellite condition, remaining useful life, and the net present value of different processing options. This flexible pricing structure incentivizes operators to utilize collection services while ensuring economic viability for service providers.

3.1.5 Novel CONOPs

The CAAS CONOPs represents a comprehensive approach to sustainable satellite constellation management that balances operational requirements with environmental responsibility through flexible infrastructure and adaptive operational procedures. Elements of the novel CAAS CONOPs can be summarized as follows:

- **Integration of Active Debris Removal with Multi-Echelon Sparing:** Combines orbital spare warehouses with ADR vehicles to create a comprehensive collection and servicing infrastructure for satellite constellation management.
- **Flexible ADR Vehicle Deployment Strategy:** Enables ADR vehicles to be deployed through multiple pathways including direct Earth launch for immediate missions, warehouse-based regional operations, or incremental fleet expansion through resupply missions.
- **Opportunistic Rendezvous Operations:** Maximizes operational efficiency by leveraging existing orbital traffic, allowing rocket stages delivering new satel-

lites to collect warehouse inventory on return trajectories and enabling Earth-launched ADR vehicles to join orbital fleets after completing collection missions.

- **Variable Collection Timing Framework:** Addresses operational urgency by requiring immediate ADR collection when necessary or immediate spare delivery when no in-plane spares are available – all while allowing flexible scheduling for in-plane spare replacement or satellite collection when time permits.
- **Multi-Pathway Satellite Processing:** Provides comprehensive end-of-life options including direct warehouse servicing and recommissioning, Earth return for refurbishment, controlled deorbiting, and terrestrial abandonment based on economic and technical considerations.
- **Incremental Infrastructure Expansion:** Enables organic growth of ADR fleet capacity through mission-based additions rather than requiring dedicated deployment launches, reducing infrastructure development costs.
- **Mandatory Collection with Operational Flexibility:** Requires ADR services for all failed satellites to ensure compliance and collision risk reduction while maintaining operational flexibility through variable timing and processing options.
- **Technology Obsolescence Integration:** Incorporates scheduled replacement of orbital spares before technological obsolescence occurs to avoid loss of revenue.
- **Dynamic Service Value Model:** Implements adaptive servicing value based on satellite condition, time to obsolescence, and net present value of processing options to incentivize collection while ensuring economic viability.
- **Warehouse Capability Scalability:** Enables incremental warehouse upgrades

to expand from basic storage + ADR refueling to full on-orbit servicing capabilities for servicing satellites.

3.2 Problem Properties and Framework Requirements

The complex, multi-domain nature of establishing circular space economies in LEO necessitates a systems-level approach that can simultaneously address technical performance, economic viability, and policy effectiveness across highly uncertain future scenarios. To answer the research question of which combination of CONOPs, flexible options, and policies consistently provide viable economic value and environmental impact of LEO-based OOS, this thesis develops a comprehensive flexibility framework that explicitly addresses the identified research gap. The methodological approach is designed to overcome the limitations of existing frameworks that fail to capture the combinatorial effects of multi-domain uncertainty, cross-fleet decision making, and the integration of CAAS.

The methodology presented here employs a three-pronged approach to properly represent the multidisciplinary nature of the research problem. First, a technical modeling component captures the operational performance of CAAS architectures and cooperative maneuvering strategies through discrete event simulation. Second, a financial analysis component evaluates economic viability using flexible option valuation methods that explicitly account for the flexibility value of various OOS strategies under uncertainty. Third, a policy modeling component incorporates the decision-making behavior of constellation operators to assess how different policy schemes influence adoption and implementation decisions.

The integration of these three components within a unified flexibility framework enables systematic screening of strategy combinations across the full range of plausible future scenarios. This approach is essential because the value proposition for LEO-based OOS emerges from the interaction between technical performance, economic

incentives, and policy environments, none of which can be meaningfully evaluated in isolation given the deep uncertainties that characterize the evolving space economy.

To establish an appropriate flexibility screening framework, we must first identify the key characteristics and properties of this problem. The framework formulation has several defining characteristics that drive the methodological choices, particularly the need for mid-fidelity methods suitable for screening designs and strategies:

1. Uncertainty is multi-variate and multi-type (both endogenous and exogenous)
2. Demand for OOS is based on several sources of uncertainty
3. The model will focus on the perspective of the satellite constellation operator, hereby referred to as the customer, assuming that the customer owns and operates the warehouses and ADR vehicles
4. Formulation should be linear and parallelizable to spare computational cost
5. Orbits are circular and maneuvers are restricted to in-plane only
6. Uncertain variables follow a pre-determined path, initialized within each Monte Carlo simulation, and is independent of decision-making
7. Decisions for the future depend on decisions that have been made in the past

3.3 Formulation Question 1: Framework Selection

Formulation Question 1

What type of flexibility framework will accommodate these properties?

In their chapter dedicated to reviewing flexibility framework literature, Cardin et al. generalize flexibility framework phases into 5 categories: baseline design, uncertainty recognition, concept generation, design space exploration, and process management [6]. As the name implies, the baseline design step determines the best design




Figure 3.7: Flexibility Framework Phases [6]

using traditional methods. In order to determine flexibility value in an infrastructure, the value of the flexible system must be compared to that of an inflexible and traditionally-designed system. The baseline design in this thesis, described further in section 3.3.1, draws inspiration from Jakob et al.'s multi-echelon sparing strategy. The uncertainty recognition phase determines which random variables should be included in the framework and how to characterize them. Section 3.3.2 provides details on the uncertain sources included within the framework. The concept generation phase decides which design elements should be included in the framework, determining which flexible options are likely to be the best candidates. The previous chapter provided justification for the proposed CAAS concept. Going into more detail, section 3.3.3 explores how this novel concept, along with other OOS CONOPs, are down-selected

and incorporated within the flexibility framework. The fourth phase, design space exploration, compares the results of the flexible designs with the baseline to determine the value of flexibility. The final phase, process management, includes consideration for the interaction of different agencies, which is out of scope for this thesis.

3.3.1 Formulation Question 1.1: Baseline Establishment

Formulation Question 1.1

How to establish the baseline?

The baseline design draws inspiration from Jakob et al.'s proposed multi-echelon sparing strategy and uses OneWeb for its use case. The experimentation introduces CAAS to the CONOPs and determine its comparative performance to a baseline configuration that utilizes overpopulation and launches ADR missions exclusively from Earth. To establish the baseline for purposes of measuring the value of flexibility, the baseline initializes all infrastructure elements and decisions from the beginning of the 30-year scenarios.

The top performing configuration parameters for the baseline design, such as the number of warehouses, their capabilities, and the refuel/repair/RPO capabilities of the satellites, are determined using the same Discrete Event Simulation and Monte Carlo method that is used within the flexibility framework, but excludes flexible options. Further details on this simulation tool are provided in section 3.4.

Parameters of the OneWeb use case are provided below in Table 3.1. While future work should consider variations to parking orbit altitude and warehouse capacity, this thesis borrows the optimal values from Jakob et al. in order to maintain a reasonable design space and focus on illustrating the value of CAAS and flexibility.

Table 3.1: Use Case Simulation Parameters

Variable Name	Range
number of planes	18
satellites per plane	36
customer altitude	1200km
parking orbit altitude	796km
inclination	86.4 degrees
number of in-plane spares	2 satellites
satellite base cost	\$900,000 [191]
satellite dry mass	150 kg[192]
satellite fuel mass	12 kg
dry mass ADR	150 kg [80]
fuel mass ADR	150 kg
ISP ADR	230 s
warehouse max capacity	35 satellites
warehouse initial capacity	5 spare satellites
warehouse fuel mass	500 kg
warehouse dry mass	1000 kg
Xenon cost	\$5000/kg [193]
Green monopropellant cost	\$100/kg [194] [195]
discount rate	0.03

Table 3.2: Sources of Uncertainty and Modeling Parameters

Uncertain Quantity	Model/Method	Initial Value(s)	Uncertainty Parameters	Eq.
Constellation Revenue	Geometric Random Walk / Log-normal PDF	\$1.4B/year [196]	$\alpha = 0.059$ [197], $\sigma = 0.15$	Equation 3.1
Launch Cost	Log-normal PDF with volatility cone	α_m based off present launch cost per kg and Citi Bear Case [198]	$\sigma = 0.1$ before 2040, $\sigma = 0.35$ and $\alpha_m = 0$ after 2040	Equation 3.1
Launch Delay	Processing + Exponential	$T_{\text{processing}} = 3$ months, $\mu_{\text{launch}} = 2$ months	Exponential with μ_{launch}	Equation 3.2
ADR Launch Delay	Learning Curve + Exponential	$\text{min}_{\text{time}} = 0.5$ years, $\text{initial}_{\text{extraTime}} = 1.5$ years, $\lambda = 0.2303$	Exponential with μ_{launch}	Equation 3.3
Satellite Manufacturing Cost	Log-normal PDF	\$900,000 [191]	$\alpha_m = -0.075$ [199], $\sigma_m = 0.1$	Equation 3.1
ADR Vehicle Cost	Learning Curve + Time-based decay	\$48M (initial) [101]	$\lambda \sim U(0.1, 0.5)$, $P_{\text{min}} \sim U(0.5, 0.8) \times \text{cost}_i$, $r = 0.03$	Equation 3.4
ADR Collection Success	Learning Curve (Power Law)	$S_0 \sim U(0.7, 0.99)$	$\lambda \sim U(0.05, 0.25)$	Equation 3.6

Continued on next page

Table 3.2 – continued from previous page

Uncertain Quantity	Model/Method	Initial Value(s)	Uncertainty Parameters	Eq.
Space-Based Services (ADR Operation, Refuel, Repair, Obsolete Repair)	Log-normal PDF	\$250k / \$250k / \$562.5k / \$687.5k	$\alpha = -0.0375, \sigma = 0.1$ (services \propto same path)	Equation 3.1
Warehouse Cost	Learning Curve	\$100M (initial)	$\lambda \sim U(0.1, 0.5), P_{\min} \sim U(0.5, 0.8) \times \text{cost}_i, N_{\max} = 54$	Equation 3.7
Warehouse Maintenance Cost	Lognormal PDF (cost)	Every 15 years: \$5M median repair cost	Cost: $\mu = \ln(5 \times 10^6), \sigma = 1.0$	Equation 3.8
Earth-Based Services (Repair, Refurbish, Obsolete Repair)	Log-normal PDF	\$450k / \$200k / \$550k (repairable); \$855k (non-repairable)	$\alpha = -0.075, \sigma = 0.1$ (services \propto same path)	Equation 3.1
Warehouse Upgrades (Refuel, Repair Capability)	Log-normal PDF	\$8M / \$15M	$\alpha = -0.0375, \sigma = 0.1$ (capabilities \propto to space-based services path)	Equation 3.1

Continued on next page

Table 3.2 – continued from previous page

Uncertain Quantity	Model/Method	Initial Value(s)	Uncertainty Parameters	Eq.
Satellite Upgrade R&D (Refuelable, Repairable)	Log-normal PDF	\$5M / \$9M R&D + 6% / 10% cost increase	$\alpha = -0.0375$, $\sigma = 0.1$ on R&D (capabilities \propto to space-based services path)	Equation 3.1
Technology Obsolescence	Weibull Utility Function	N/A	$k = 2$, $\lambda = 1$, $\beta = 2$	Equation 3.9
Failure Time	Exponential (from MTBF) and Failure Change Rate, Uniform	MTBF starting with 1 fail- ure/year, includes exponen- tial change rate with random parameter from uniform dis- tribution	MTBF Derived from fleet size, exponential failure change rate varies uni- formly between 0.005 and 0.08	Equation 3.10, Equation 3.11, and Equation 3.12
Failure Type	Bernoulli Trial	50% inoperable	None (uniform draw)	—
Collision Event Cost	Exponential + Uniform + Cascade Probability	\$10k base	$\beta = 0.01$, $P_{\text{cascade}} = \min(0.02T, 0.5)$	Equation 3.13

3.3.2 Formulation Question 1.2: Uncertainty Variable Selection and Characterization

Formulation Question 1.2

How to identify and characterize sources of uncertainty?

There are a number of uncertain variables involved in a space mission. This list includes, but is not limited to: launch cost, spacecraft cost, schedule delays, random failures, change in demand, industry growth, technology obsolescence, technology readiness, and learning curves [134] [142] [3]. This framework focuses on mega-constellation revenue, technology obsolescence, launch cost, random failures, and the hybrid uncertainty arising from interactions between servicing capability and satellite capability. Table 3.2 summarizes sources of uncertainty and their characterization while Table 3.3 contains the uncertain variables that were randomized using a DOE. This section provides a deep dive into the details of each source of uncertainty and its characterization.

Table 3.3: Latin Hypercube DOE Parameters for Uncertain Variables

Parameter	Min	Max
Earth-based repair operation cost streamlining cost multiplier	0.5	1
Earth-based refuel operation streamlining cost multiplier	0.5	1
Technology obsolescence intensity parameter	1	5
Time to technology obsolescence onset (years)	5	15
Active Debris Removal (ADR) initial success rate	0.7	0.99
ADR success rate learning curve exponent	0.05	0.25
Warehouse cost learning curve exponent	0.1	0.5
Warehouse cost minimum cost fraction	0.5	0.8
ADR cost learning curve exponent	0.1	0.5
ADR minimum cost fraction	0.5	0.8
Satellite cost multiplier for RPO-capability upgrade	1	1.5

Parameter	Min	Max
Satellite cost multiplier for repair upgrade	1	1.3
Satellite cost multiplier for refuel upgrade	1	1.2
Satellite return to Earth cost (multiplier of present launch cost)	0.2	0.99
Simulation year that returning satellites to Earth becomes possible	7	15
Satellite failure rate change rate (per year)	0.005	0.079
Percentage of launch vehicles with reusable 2nd stages	20%	70%
Reusable 2nd stage fleet percent multiplier if Earth-return available	1 \times	1.25 \times

Geometric Random Walk for Revenue Modeling

Revenue uncertainty and random failure modeling are commonly included in related flexibility frameworks [10] [20]. The dramatic reduction in launch costs has fundamentally transformed the commercial space industry over recent decades, making continued launch cost uncertainty a critical consideration. While current mega-constellation operators plan continued expansion, their growth trajectories remain uncertain and are directly tied to customer revenues and profit margins.

The geometric random walk method provides a robust foundation for modeling revenue uncertainty and has been extensively used in previous on-orbit servicing (OOS) flexibility frameworks [10, 141, 142]. This approach yields the following log-normal probability density function:

$$p_{\tau}^{(m)}(x) = \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma_m \sqrt{\tau}} \frac{1}{x} \exp \left\{ -\frac{(\ln(x) - (\alpha_m - \sigma_m^2/2) \tau)^2}{2\sigma_m^2 \tau} \right\} \quad (3.1)$$

In this formulation, α_m represents the drift (1/time), σ_m denotes volatility (1/ $\sqrt{\text{time}}$), and τ represents time. The drift parameter, α_m , captures the time value of money [10]. This geometric random walk model is applied to various sources of uncertainty throughout the framework, providing a consistent stochastic foundation for financial

modeling.

Launch Cost and Delay Uncertainty

Launch cost scenarios utilize log-normal models incorporating volatility cones based on 2020 Citi Research projections [198]. By 2040, Citi Research predicts Falcon Heavy launch costs will range from a bull case of \$33 per kg to a bear case of \$300 per kg, with a base case of \$100 per kg.




Figure 3.8: Launch Cost Predictions [198]

Using a similar log-normal relationship with cone of uncertainty, as demonstrated by Lamassoure with 30% volatility applied to LEO constellation market forecasts [142], piece-wise linear representations of the Citi Research launch cost projections provide probabilistic characterization of this critical uncertain parameter. We can assess the accuracy of these projections by comparing their predictions for 2020 to 2025 with the actual present-day launch costs: launch cost for Falcon 9 is presently \$3,986 [200], not \$400. This thesis therefore uses a simplified launch cost projection model with an α_m based off the linear curve between present cost (\$3,986, 2025) and the Citi bear case in 2040 (\$300, 2040). After 2040, $\alpha_m = 0$ to represent the launch cost leveling out. The volatility parameter σ_m is approximated at 0.1 until 2024.

Since the launch cost projections only go until 2040, σ_m increases to 0.35 after year 2040 to reflect increased uncertainty.

Launch timing uncertainty follows an exponential distribution model, with processing time $T_{\text{processing}} = 3$ months and mean launch delay $\mu_{\text{launch}} = 2$ months:

$$T_{\text{launch}} = T_{\text{processing}} + X_{\text{launch}}, \quad X_{\text{launch}} \sim \text{Exponential}(\mu_{\text{launch}}) \quad (3.2)$$

For Active Debris Removal (ADR) missions launched directly from Earth, processing time follows a learning curve as manufacturing and launch experience accumulates:

$$\text{adrTprocessing} = \min_{\text{time}} + (\text{initial}_{\text{extraTime}} \times e^{-\lambda t}) \quad (3.3)$$

where $\min_{\text{time}} = 0.5$ years, $\text{initial}_{\text{extraTime}} = 1.5$ years, and $\lambda = 0.2303$.




Figure 3.9: Cone of Uncertainty, Geometric Diffusion [142]

Additionally, this framework assumes that downmass payload capacity on reusable

second stages is 1/3 of its launch payload capacity.

ADR Vehicle Cost Modeling

ADR vehicle costs are modeled using a time-based exponential decay starting from an initial cost of \$48,000,000 [101]. The cost incorporates a learning curve effect based on the number of vehicles manufactured. Equation 3.4 provides the cost multiplier

$$n = \{1, 2, \dots, N_{\max}\}; \text{value}(n) = P_{\min} + (1 - P_{\min}) \cdot e^{-\lambda n} \quad (3.4)$$

where n represents the number of ADR vehicles added (starting from 1), N_{\max} is the maximum number of ADR vehicles considered (assumed to be equal to the number of planes, 18), P_{\min} is the minimum ADR cost fraction (asymptotic limit), which is randomly sampled between 0.5 and 0.8, and λ is the learning exponent controlling the rate of decay, which is randomly sampled between 0.1 and 0.5.

Additionally, a time-based cost reduction applies with annual reduction rate $r = 0.03$ and time t in years:

$$\text{cost_reduction} = (1 - r)^t \quad (3.5)$$

Space-Based Operations and Costs

ADR vehicles experience variable collection success rates that improve according to a learning curve trend:

$$S_n = 1 - (1 - S_0) \cdot n^{-\lambda}, \quad \forall n \in \{1, 2, \dots, N_{\max}\} \quad (3.6)$$

where S_n represents the probability of successfully collecting the n -th satellite, n is the number of collected satellites (starting from 1), N_{\max} is the maximum number of collected satellites considered to impact the learning curve (assumed to be 50,

after which the collection probability remains constant), S_0 is the initial collection probability ranging between 0.7 and 0.99, and λ is the learning exponent controlling the improvement rate, ranging between 0.05 and 0.25.

Space-based operation costs begin at \$250,000 per operation and evolve following the same geometric random path as other space-based services, as these capabilities are assumed to be interrelated. The ADR operation cost for RPO satellites is one third the typical ADR operation cost (\$83,250 per operation) because RPO satellites are upgraded to be cooperative targets, and the maneuver duration is shorter due to this enhanced capability.

Warehouse Cost Uncertainty

Warehouse cost uncertainty mirrors the ADR vehicle cost modeling approach, with a total cost of \$100,000,000 that decreases with each new warehouse added to the system. This cost includes the manufacturing and operation cost, but does not include launch cost or the cost of fuel, spare, and upgrade resupplies. The learning-curve-based reduction follows:

$$n = \{1, 2, \dots, N_{\max}\}; \text{value}(n) = P_{\min} + (1 - P_{\min}) \cdot e^{-\lambda n} \quad (3.7)$$

where n is the number of warehouses added (starting from 1), N_{\max} is the maximum number of warehouses considered, P_{\min} is the minimum performance level (asymptotic limit), and λ is the learning exponent (rate of decay) randomly sampled between 0.1 and 0.5.

Warehouse Maintenance Cost Modeling

In addition to initial construction and operational costs, warehouses require maintenance every 15 years. The cost associated with each warehouse repair is modeled using a lognormal distribution to capture the high variability and right-skewed nature

of major infrastructure repair costs:

$$C_{\text{failure}} \sim \text{Lognormal}(\mu = \ln(5 \times 10^6), \sigma = 1.0) \quad (3.8)$$

where the median failure cost is approximately \$5 million, while $\sigma=1.0$ allows for significant variability. This modeling approach reflects the reality that warehouse infrastructure failures can range from minor system repairs to complete facility reconstruction, depending on the nature and severity of the failure event.

The total warehouse-related costs over the mission lifetime thus include initial construction costs (with learning curve and time-based reductions), operational costs, and stochastic failure repair costs that occur throughout the warehouse operational life.

Technology Obsolescence Modeling

Technology obsolescence represents another critical uncertainty source, previously recommended for future work by Lamassoure [142]. Since then, researchers have developed sophisticated technology obsolescence models for spacecraft [201]. Given the rapid pace of advancement in the commercial space industry, technology obsolescence has become increasingly important, with many spacecraft operators choosing LEO and shorter life cycles specifically for the flexibility to upgrade satellites as technology evolves.

Technology obsolescence is captured using a Weibull-based utility function that reduces satellite revenue after reaching obsolescence time, based on the utility function $u(t)$ developed by Geng et al. [201]. The model employs a 3-parameter Weibull distribution with $\beta = 2$. The intensity metric is randomly determined for each scenario, with time to obsolescence drawn from a Weibull distribution having shape parameter $k = 2$ and scale parameter $\lambda = 1$. The index, i , represents the individual revenue contributions of each satellite in the constellation.

$$u_i(t) = u_{o,i} e^{-\left(\frac{(t-T_{obs,i})}{\theta_{obs,i}}\right)^\beta} \quad \text{for } t \geq T_{obs,i}, \quad (3.9)$$

$$u_{\text{total}} = \sum_i u_i(t).$$

Geng et al. demonstrate how technology obsolescence impacts spacecraft utility degradation and Net Present Value (NPV) calculations [201]. This framework applies this approach to customer NPV calculations to capture rapidly evolving satellite technology effects. When customers opt to receive servicing, the obsolescence timer T_{obs} resets, providing a mechanism for technology refresh. In each scenario, the time to obsolescence (T_{obs} in years) and intensity of obsolescence (θ_{obs} in years), are randomly selected from ranges [7,20] and [1,5], respectively.




Figure 3.10: Spacecraft Present Value with Technology Obsolescence [201]

Satellite Failure Rate Modeling

Based on OneWeb's observed experience of 4 satellite failures over 4 years [202] [203], this framework sets the baseline failure rate to 1 satellite failure per year. Mean Time Between Failure (MTBF) is calculated as:

$$\text{MTBF} = \left(\frac{1}{\text{failureRate}/(\text{numPlanes} \times \text{totalSatellitesPerPlane})} \right) \quad (3.10)$$

Each satellite receives an assigned failure time based on the MTBF using an exponential distribution:

$$T_{\text{failure}} \sim \text{Exponential} \left(\frac{1}{\text{MTBF}} \right) \quad (3.11)$$

While simulations begin with a baseline of 1 failure per year, this failure rate decreases over time to reflect improving satellite reliability:

$$f = f_0 - \min(0.1 \cdot e^{r \cdot t}, 0.99) \quad (3.12)$$

where f is the current failure rate, f_0 is the original failure rate, r is the failure rate change rate (randomly sampled from a uniform distribution between 0.005 and 0.08), and t is current time.

This approach differs from previous work by Luu et al., who used a Markov process with failure rate 0.3 approximated from the Iridium constellation [3], and Saleh et al., who set failure rate to $\lambda = 1/\text{mean mission length}$ [141]. Jakob et al. approximate OneWeb and Starlink constellation failure rates at 0.05 failures per year [14]. The framework in this thesis, however, uses more recent observational data and includes temporal improvement trends.

Having defined failure and obsolescence properties, it is important to distinguish the three different satellite states that are modeled in this framework: retirement, failure, and obsolescence. When a satellite is retired, it means that it has run out of propellant aside from its designated deorbit fuel and requires refueling service in order to extend its lifetime. When a satellite is obsolete, it provides reduced revenue if it is active in the constellation, contributing to indirect cost. Once it is repaired,

either on Earth or in a capable warehouse, its time to obsolescence is reset. When a satellite fails, it requires repair in order to have any functionality again.

Operational Uncertainty and Failure Classification The simulation framework addresses operational uncertainties, with particular emphasis on collision risks arising from satellite failures. The simulation distinguishes between failure types, where certain failures result in completely inoperable satellites. Constellation operators face regulatory requirements to secure new FCC authorization for replacement satellites. The FCC evaluates each application considering the cumulative collision risk presented by the constellation [204][2].

Active debris removal services function as an insurance mechanism for constellation operators, preventing collision-related costs that could substantially exceed the ADR mission expenses. Based on these considerations, the framework operates under the assumption that all failed satellites receive ADR services, with 50% of failed satellites experiencing replacement authorization delays until the failed satellite is removed.

Collision Avoidance and Collision Event Costs

Collision probability depends on the number of failed satellites and mission duration:

$$P_{\text{collision}} = 1 - e^{-\beta \cdot N_{\text{fail}} \cdot T} \quad (3.13)$$

where N_{fail} is the number of failed satellites, T is the number of years, and $\beta = 0.01$ represents the collision hazard rate per satellite per year. The base cost of collision-related events is defined as $C_{\text{base}} = \$10,000$.

The probability of cascade-type catastrophic events increases with time:

$$P_{\text{cascade}} = \min(0.02 \cdot T_{\text{cascade}}, 0.5) \quad (3.14)$$

Collision costs are drawn from uniform distributions that depend on whether a cascade event occurs:

$$C_{\text{collision}} = \begin{cases} C_{\text{base}} \cdot U(1000, 10000), & \text{with probability } P_{\text{cascade}} \\ C_{\text{base}} \cdot U(1, 1000), & \text{with probability } 1 - P_{\text{cascade}} \end{cases} \quad (3.15)$$

The actual incurred cost is:

$$C = \begin{cases} C_{\text{collision}}, & \text{if a collision occurs (with probability } P_{\text{collision}}) \\ 0, & \text{otherwise} \end{cases} \quad (3.16)$$

Operational and Cost Uncertainties

Beyond the uncertainty sources characterized above, the CAAS framework incorporates numerous operational and cost parameters that exhibit significant epistemic uncertainty. These parameters represent aspects of the system where limited historical data, emerging technologies, or novel operational concepts prevent confident point estimates. Rather than making potentially erroneous assumptions about these uncertain parameters, this framework employs Latin Hypercube Sampling (LHS) within a Design of Experiments (DOE) approach to systematically explore the parameter space of these critical unknowns.

Latin Hypercube Sampling provides an efficient method for exploring high-dimensional parameter spaces while ensuring representative coverage across all parameter ranges. LHS divides each parameter's range into equal probability intervals and samples exactly once from each interval, ensuring more uniform exploration of the parameter space with fewer samples. This approach is particularly valuable for computationally

intensive simulations where exhaustive parameter exploration would be prohibitively expensive.

The operational and cost uncertainties incorporated through Latin Hypercube DOE are summarized in Table 3.3. These parameters represent key operational aspects and cost drivers where data is sparse, emerging technology performance is uncertain, or novel operational concepts lack empirical validation.

Operational Efficiency and Learning Parameters Several parameters capture the potential for operational efficiency improvements through learning and scale effects. Earth-based repair and refuel operation cost streamlining multipliers (0.5-1.0) acknowledge that ground-based servicing operations may benefit from economies of scale and satellites that are upgraded to be refuelable and repairable in space. The lower bound of 0.5 represents a scenario where operational efficiencies reduce costs by 50%, while the upper bound of 1.0 assumes no efficiency gains beyond baseline estimates.

Learning curve exponents for both warehouse costs (0.1-0.5) and ADR costs (0.1-0.5) reflect uncertainty in how quickly experience accumulation translates to cost reductions. Lower exponent values indicate slower learning, while higher values suggest rapid cost improvements with experience. The corresponding minimum cost fractions (0.5-0.8 for both warehouse and ADR systems) represent the asymptotic cost level achievable through learning, acknowledging that some baseline costs cannot be eliminated regardless of experience.

ADR Technology Performance and Maturation ADR operational performance parameters acknowledge the nascent state of active debris removal technology. Initial success rates (0.7-0.99) span from moderately reliable early operations to near-perfect performance as systems mature. The ADR success rate learning curve exponent (0.05-0.25) captures uncertainty in how quickly operational experience translates to

improved collection reliability.

Satellite Upgrade and Enhancement Costs Satellite upgrade cost multipliers reflect the uncertain costs of incorporating additional capabilities into satellite designs. RPO (Rendezvous and Proximity Operations) capability upgrades (1.0-1.5x baseline cost) represent the most significant design modifications, requiring additional sensors, computing power, and maneuvering capability. Repair upgrade costs (1.0-1.3x) assume moderate increases for standardized interfaces and diagnostic capabilities. Refuel upgrade costs (1.0-1.2x) represent the smallest premium, reflecting relatively mature refueling interface technologies. It's important to note that while this framework includes cost penalties related to satellite upgrades, it assumes satellite upgrades to do incur significant mass penalties.

These ranges recognize that satellite upgrade costs depend heavily on implementation approach, integration complexity, and economies of scale across constellation deployment. Early implementations may incur significant cost premiums, while mature designs with standardized interfaces may approach baseline satellite costs.

Novel Operational Concepts Satellite return to Earth operations represent a novel concept with substantial cost uncertainty. The cost multiplier relative to present launch costs (0.2-0.99) spans from highly efficient return systems that cost only 20% of launch costs to expensive return operations approaching full launch costs. This wide range reflects fundamental uncertainty about return system architectures, reusability, and operational efficiency.

The timing parameter for when satellite return becomes feasible (years 1-15) acknowledges uncertainty in technology development timelines and regulatory approval processes. Early availability (year 7) assumes rapid deployment of return capabilities, while later availability (year 15) reflects more conservative technology development and certification timelines.

System Reliability Evolution The satellite failure rate change rate (0.005-0.079 per year) captures uncertainty in how satellite reliability improves over time. This parameter affects the exponential improvement model for satellite failure rates, with lower values representing gradual reliability improvements and higher values indicating rapid advances in satellite design and manufacturing quality.

Justification for Parameter Ranges The ranges are intentionally broad to capture genuine epistemic uncertainty while remaining physically and economically reasonable. Conservative bounds ensure that extreme scenarios remain plausible rather than representing purely academic exercises.

This approach explicitly acknowledges uncertainty rather than masking it with arbitrary point estimates. By systematically exploring these parameter ranges through Latin Hypercube sampling, the framework provides insights into which uncertainties most significantly impact system performance, enabling targeted research and development efforts to reduce critical uncertainties.

Some costs have a greater basis for estimation and are incorporated as point-estimates within this framework. These cost assumptions and their justification are presented in Table 3.5. However, there is still uncertainty regarding these cost assumptions, so sensitivity testing on select deterministic values is conducted to determine if perturbations impact overall rank order of top performing configurations. Details on this portion of the experimentation is provided in subsubsection 4.3.3.

Uncertainty Incorporation and Monte Carlo Implementation

The framework generates Monte Carlo scenarios where each scenario contains pre-determined paths of uncertain variables sampled from their respective distributions. These sources include mega-constellation revenue, technology obsolescence, launch cost, collision events, and random failures, all assumed independent of customer

decision-making. Each Monte Carlo scenario creates a 30-year timeline for these uncertain variables, initialized at the start of the simulation.

For this screening framework, assumptions and simplifications preserve the simulated effect of uncertainty while managing computational cost and model complexity.




Figure 3.11: Monte Carlo Scenario Construction

This comprehensive uncertainty characterization enables the flexibility framework to evaluate Collection-as-a-Service architectures under realistic uncertain conditions, supporting robust decision-making for sustainable space operations and the advancement of a circular space economy.

3.3.3 Formulation Question 1.3: Concept Generation

Formulation Question 1.3

How to develop the concept generation?

The concept generation phase determines which design elements should be in-

cluded in the framework, identifying the most promising flexible mechanisms and strategies. The literature review, described in the previous chapter, revealed critical limitations in previously proposed LEO OOS concepts that informed the concept generation decisions of this thesis. Proliferated servicing pods, while innovative, demonstrated insufficient value proposition, providing benefits only in narrow scenarios with high failure rates when timeliness received priority weighting. Similarly, traditional one-to-one and one-to-many servicing CONOPs fails to address the fundamental economic and maneuvering challenges of LEO operations, particularly the cost-effective nature of proliferated constellations and the dominance of established sparing strategies. Cooperative maneuvering showed promise but remained limited to chemical propulsion systems and one-to-many architectures, failing to address the electric propulsion systems used by most mega-constellations.

The concept generation phase strategically focused on three key design elements that address these identified gaps. CAAS emerged as the primary concept because it builds directly upon Jakob et al.'s multi-echelon sparing strategy, which is an approach that already appeals to satellite operators' fundamental needs for cost reduction and coverage maintenance. Rather than competing with spare strategies, CAAS extends and enhances them by transforming spare warehouses into flexible collection hubs that can upgrade to provide servicing capabilities. This approach addresses the classic "chicken-and-egg" problem through modularity and incremental upgrades, allowing the infrastructure to evolve without requiring massive upfront investment in unproven technologies.

Cooperative maneuvering was retained but re-conceptualized beyond the scenarios previously studied. The framework incorporates cooperative maneuvers within the context of many-to-many servicing architectures and approximates the cost and duration associated with electric propulsion systems, making it more applicable to contemporary mega-constellation operations. Servicing at designated warehouses was

prioritized over proliferated pods because warehouse-based operations provide superior economies of scale, leverage J2 drift effects for efficient plane changes, and offer the flexibility to incrementally add capabilities as market conditions evolve.

These concept selections collectively address the core flexibility requirements identified in the literature: providing immediate value to customers through enhanced sparing capabilities, creating pathways for incremental technology adoption, enabling economies of scale through centralized operations, and maintaining compatibility with existing operational paradigms. By building upon proven sparing strategies rather than replacing them, the selected concepts reduce implementation risk while preserving optionality for future technological developments, embodying the fundamental principles of flexible system design under deep uncertainty.

3.3.4 Formulation Question 1.4: Design Space Exploration

Formulation Question 1.4

How to conduct design space exploration?

Methods for solving real options models include binomial trees, finite difference methods, and Monte Carlo simulations, to name a few. While binomial trees and related methods are computationally efficient, they are only useful for problems with a single source of uncertainty. Binomial lattices and decision trees suffer from the curse of dimensionality for such a large design space. Simultaneously including multiple sources and multiple flexible options requires multinomial lattices, which complicate both decision policies and computation [205]. Additionally, binomial lattices rely on path-independence, which is a valid assumption for option pricing models when markets are efficient, but is not always valid for most engineering systems. As Cardin et al. note, engineering systems are sensitive to initial conditions, inputs, and past events, so it is important to be mindful of past history when making a decision

about the future [205]. Simulation-based methods not only handle multiple sources of uncertainty, they also offer more flexibility for uncertain parameters' probability distributions.

Decision analysis employs decision trees and backward induction (as used in dynamic programming, [206]), which is often intractable for frameworks with multiple options and multiple uncertain variables. Discretized scenarios, often paired with stochastic optimization, use either direct formulation or decision-rule formulation to identify flexible designs [207]. Hamdan et. al utilize staged decision analysis for complex, stochastic, and time-variant systems and capture the effect of evolving uncertainty.

Stochastic programming often requires advanced methods, therefore, they are often represented with deterministic equivalent formulations [207]. Meanwhile, simulations paired with decision rules can account for multiple options and uncertain variables via scenarios [6]. The simulation method can be computationally expensive due to the number of scenarios and decision rule variables, so it is useful to simplify simulation models such that runtime is acceptable and system behavior is still properly represented.

Formulation Question 1.4.1: Analogous Frameworks

Formulation Question 1.4.1

Are there any analogous flexibility frameworks with similar properties?

A 2013 offshore oil infrastructure framework, previously introduced in section 2.9.2, paired a bottom-up, integrated systems model with Monte Carlo simulations and decision rules to screen the value of flexible options [12]. This framework contains many of the same desired characteristics for the proposed OOS flexibility framework in this thesis, such as:

1. Multi-level flexibility
2. Large design space
3. Multiple sources of uncertainty
4. Monte Carlo scenarios
5. Combinatorial effect of options

Observing that recent oil and gas basin discoveries are small and dispersed, Lin et al. identify subsea tiebacks that connect new basins to existing offshore facilities as a means of leveraging these discoveries without needing to build new, expensive rigs [208]. This is especially useful given the volatility in the oil and gas industry as well as uncertainties about basin size and facility performance. There are many challenges regarding the design and operations of subsea tiebacks, such as where to place them and when. There is also the question of where to place new facilities and whether these facilities should have the ability to expand capacity in the future.

This application is analogous to OOS for proliferated constellations in LEO because it considers flexible options to leverage oil from small and scattered sources rather than from a single, large, accessible oil field, of which there are fewer and fewer. Monolithic oil fields are akin to large satellites in GEO, which provide a more immediate business case for supporting infrastructure. With flexibility and economies of scale, however, novel, flexible enablers such as subsea tiebacks and collection hubs can pave the way for innovative infrastructures that better utilize available resources.

After identifying multi-level flexibilities, such as the various combinations of fields and facilities, Lin et al. develop an integrated systems model to represent the project life time and determine its net present value (NPV). Flexible options are implemented when conditions trigger their associated decision rules. By running multi-variate Monte Carlo scenarios in both flexible infrastructures and inflexible baseline

models, Lin et al. are able to examine the best options, and combination of options, by conducting statistical analysis. The outer Monte Carlo loops creates scenarios that characterize the state of infrastructure elements and the market in one instance. Meanwhile, the inner loop simulates the entire lifecycle of the oil field given the uncertainties sampled within the outer loop. The inner loop triggers the decision rules that are responsible for infrastructure reconfiguration. Embedded loops captures uncertainty evolution over the course of the project.

In their 2008 thesis, Lin develops this screening framework and provides further explanation for the methodology choices. Lin determines that mid-fidelity models are sufficient for screening options because they represent system behavior while providing quick computational time. High-fidelity models at the early conceptual design phase obfuscate the big picture and can potentially mislead the user, given the lack of complete and accurate inputs [208]. Meanwhile, low-fidelity models could oversimplify the dynamics of a complex system. Lin's mid-fidelity integrated screening models represent the feedback loop between inputs, production systems, and output systems, capturing the physical flows, logical flows, and financial flows between elements. Mid-fidelity models take on the order of seconds to minutes to run and provide prediction error less than 10 to 20%. For this reason, mid-fidelity models are suitable for screening design options ahead of detailed design.

Lin considers a taxonomy of different flexibility levels within their framework. Strategic flexibility involves high-level technological concepts and configurations. Tactical flexibility is the system's ability to modify behavior and performance, but not the overall architecture. Lastly, operational flexibility means that a system can adjust its operations to meet present conditions [208].

To trigger the implementation of flexible options, Lin utilizes conditional decision rules, embedded within decision trees. When certain conditions are met, the framework implements the given flexible options. Lin tunes the decision trees for

each flexible option with sensitivity studies. These decision rules are parametrically tuned, but not optimized. However, this approach still provides sufficient performance improvement when compared to the rigid baseline system.

Overall, Lin's thesis presents a four-step process for screening flexible options under multiple sources of uncertainty, containing a modeling, strategy synthesis, simulation, and screening phase. An important part of the process is using the framework as a computational laboratory to determine the best decision parameters and best combination of flexible options. Lin et al.'s experimentation process, which is revisited in chapter 4, uses various designs of experiments (DOEs) to explore the design space.

They measure infrastructure performance over the various Monte Carlo simulations by comparing value-at-risk/gain (VARG) plots of the results which provides the cumulative distribution of Net Present Value (NPV). An example of VARG curves, shown in Figure 3.12, illustrate how flexible options can influence its shape. If the VARG plot depicts profit, option B, compared to option A, provides higher value at gain but lower value at risk and a lower average value. If curve A corresponds to a flexible configuration, it illustrates how to added flexibility increases average profits and alleviates risk at the 5th percentile compared to the inflexible configuration, curve B. The switch from the inflexible configuration (curve B) to the flexible configuration (curve A) does come with a trade-off, however, since the flexible option doesn't provide the same opportunity for large profits that B provides.

For their screening framework, Lin et al. produce a series of Monte Carlos that sample from uncertainty distributions to create evolving trajectories [208]. When the trajectory experiences a discrete jump, the range of uncertainty shifts with it. The first step of their process is to create an initial distribution vector. Then, they update the distribution vector at discrete time steps, applying random walk with decreasing variation over time. As described in section 3.3.2, this thesis takes a simpler approach




Figure 3.12: Example of a VARG Plot (Profit)

to characterizing uncertainty, creating uncertain scenarios at the initialization of each simulation that are not impacted by decisions.

In summary, to address the identified requirements and gaps, the flexibility framework presented in this thesis will screen flexible options and assess the impact of policy by using simulation-based methods with decision rules and Monte Carlo scenarios. This is a departure from previous OOS flexibility frameworks, which used classic real options analysis with the backward induction process [142] [10] [141]. Details on decision rule implementation are explained in section 3.5. All decision rule algorithms are provided in Appendix B.

3.4 Formulation Question 2: Simulation Method Selection and Development

Formulation Question 2

What simulation method is appropriate?

3.4.1 Space Logistics Modeling

Having established the structure of the flexibility framework, it is necessary to determine the appropriate simulation method to capture the described CONOPs and decision-making process. The simulation must be a computationally-efficient space logistics model that manages the logistics of the OOS infrastructures while maintaining mid-level fidelity.

Network flow modeling stands as the state-of-the-art approach for space logistics, having been perfected and tailored over a wide range of applications. Originating from graph theory, network flow offers a distinct advantage due to its compatibility with Mixed Integer Linear Programming (MILP), which enables efficient attainment of global optima [93]. Space logistics requires a number of adjustments to traditional network flow methods, starting with a generalization to include flows with gains and losses [209]. Additionally, space logistics networks frequently contain several resources that flow across arcs, necessitating the multi-commodity network flow formulation. When commodities interact and flow is not conserved along arcs, it is necessary to use a generalized multi-commodity network flow (GMCNF) [210] [211]. GMCNF problems are only capable of representing static networks, however, so for problems with time-dependent characteristics, it is useful to use a time-expanded network approach. Ho establishes a dynamic network framework that incorporates evolving technology options and introduces time-expanded networks with linear programming [212] [213].

Since the development of these network flow formulations, there has been an expansive body of work in the space logistics field that utilizes Time-Expanded, Generalized Multi-Commodity Network Flow, Mixed Integer Linear Programming, and their variations [214] [215] [215] [216] [217]. Tristan Sarton du Jonchay's 2022 dissertation, entitled "Framework for the design and operations of sustainable on-orbit servicing infrastructures dedicated to geosynchronous satellites," provides a framework that systematically determines optimal routes for on-orbit servicers in GEO

[20]. This framework is an adapted traveling salesman problem, using the Time-Expanded, Generalized Multi-Commodity Network Flow (TE-GMCNF) method to capture time-sensitive, multi-arc behavior between nodes. The traveling salesman method alone is insufficient for on-orbit servicing networks, because it fails to incorporate demand uncertainty and commodity flow to the service depots. Sarton du Jonchay uses rolling horizons to adjust the logistics scheduling as changes in demand (random failures) occur. At each horizon, Sarton du Jonchay implements mixed-integer linear programming (MILP) to optimize logistics scheduling.

In Sarton du Jonchay's operational concept, the OOS company launches and manages orbital depots and impulsive-thrust service vehicles to service GEO satellites on both a random and planned basis. The OOS company decides whether to provide service when a potential customer experiences a random failure. If they decide to proceed, they send a servicer. Failure to meet the servicing deadline results in monetary penalty. Following service, the dispatched service vehicle is available for its next operation. Scheduling and supply chain are optimized concurrently. Sarton du Jonchay includes a number of service options, such as inspection, refueling, station-keeping, satellite repositioning, repair, mechanism deployment, and retirement. Resources, such as vehicles, spares, and fuel, are regularly supplied from Earth via launch vehicles. Service vehicles include notional tools such as refueling apparatuses, observation sensors, dexterous robotic arms, and capture mechanisms. Servicers utilize fully autonomous or semi-autonomous robotics. Sarton du Jonchay uses the framework to compare architectures with specialized servicers (containing only the tools to do specific tasks) and generalized servicers (capable of doing many tasks). There are three journal publications related to the thesis. The first paper assumes customer satellites are distributed in the same circular orbit and only uses high-thrust equations [93]. The second paper includes both low-thrust and high-thrust maneuvers in the same circular orbit [218]. The third paper allows for multi-orbit maneuvering [219].

Like Lamassouire, Sarton du Jonchay’s demand uncertainty is based on random services. The framework doesn’t include policy or environmental considerations, nor does it include options for long-term collection as an on-orbit servicing option. Like many other OOS frameworks and methodologies that argue GEO provides the best OOS business case, Sarton du Jonchay’s thesis focuses exclusively on GEO. Sarton du Jonchay’s framework does not include incremental deployment or an assessment of the value of flexibility, but it could be adapted to do so. They note that the framework could accommodate a dynamic market forecast and provide insight on useful infrastructure changes and updates when the framework is subject to sensitivity analysis [20].

3.4.2 Discrete Event Simulation

While state-of-the-art, network flow models with optimization provide a higher degree of fidelity than is necessary for a screening framework designed to evaluate flexible options. The fundamental requirements established by Lin’s flexibility framework [12] necessitate a simulation approach that can efficiently handle: (1) multi-level flexibility with numerous option combinations, (2) multiple sources of uncertainty through Monte Carlo scenarios, (3) dynamic decision rules that respond to evolving conditions, (4) computational tractability for extensive scenario analysis, and (5) sufficient fidelity to distinguish between alternative strategies without unnecessary complexity.

While Sarton du Jonchay’s formulation for TE-GMCNF with Rolling Horizons matches several characteristics of the problem presented in this thesis, the degree of necessary adjustments and additional constraints would prove not only computationally intractable, but it would also provide a higher degree of fidelity than is necessary for the purposes of a flexibility framework that seeks to compare options to one another. The complexity of incorporating the CAAS CONOPs, including multi-level

decision-making, incremental infrastructure deployment, and dynamic upgrading of warehouses, would require extensive modifications to existing MILP formulations, significantly increasing computational burden without proportional gains in insight for option screening purposes.

Discrete Event Simulation emerges as the appropriate methodological choice for several key reasons that align with Lin’s framework requirements. First, DES naturally accommodates the time-dependent, event-driven nature of satellite operations, failures, and servicing decisions without requiring the discretization and linearization constraints inherent in network flow approaches. Second, DES enables the straightforward implementation of complex decision rules that can adapt to evolving system states and uncertainty realizations. Third, the object-oriented nature of DES frameworks allows for modular representation of system components (satellites, ADR vehicles, warehouses) with individual properties and behaviors that can be easily modified to test different flexible options.

DES has been successfully employed in space logistics applications, demonstrating its suitability for this domain. Sears et al. use DES to investigate the impact of on-orbit recycling and manufacturing capabilities on the value of on-orbit servicing [86], while DES is also frequently employed in analogous logistics applications, such as maintenance strategies for the civil aviation industry. The computational efficiency of DES enables the extensive Monte Carlo simulation required to explore the uncertainty space defined in Lin’s framework, allowing for robust statistical analysis of flexible option performance across thousands of scenarios.

A discrete event simulation provides a sufficient test bed to compare the value of options while maintaining the computational tractability necessary for comprehensive uncertainty analysis. The DES approach implemented in this research utilizes event-driven time advancement, state-dependent decision rules, and continuous monitoring of system performance metrics that enable the identification of conditions under which

flexible options provide value.

The DES framework's ability to incorporate stochastic elements, such as random satellite failures, uncertain demand patterns, and technology obsolescence, through Monte Carlo sampling aligns perfectly with Lin's emphasis on scenario-based uncertainty representation. Moreover, DES naturally accommodates the hybrid uncertainties arising from interactions between servicing-infrastructure-level decisions, satellite-fleet-level decisions, and individual-satellite-level decisions, which are central to the CAAS value proposition but would be challenging to represent in optimization-based approaches without significantly increasing problem complexity.

3.5 Formulation Question 3: Decision Rule Creation and Calibration

Formulation Question 3

How to create and calibrate decision rules?

Decision rules are commonly featured in flexibility frameworks. Chen et al. pair decision rules with multi-stage stochastic programming to mitigate against launch delays for ISS resupply missions [134]. Their decision rules and Pareto front of expected mission cost vs expected mission performance allow the user to determine how many resources to launch and how to avoid lost productivity. Chen et al. note that decision rules don't require gradients and can be directly incorporated into mixed-integer linear programming formulations. Cardin et al. also use decision rules in their multi-stage stochastic programming framework, finding that optimization with decision rules provide a sufficient approximation of optimal solutions, comparing their results with those found by Real Option Analysis (ROA) [205]. In a different paper, Cardin et al. apply the decision rule approach to on-demand vehicle-sharing [136].

The OOS flexibility framework developed by Saleh, Lamassoure, and Hastings uses a form of decision tree analysis to trigger flexible options at each decision point [141]. This decision is made from the customer's perspective and reacts to the present

operational mode and present state of uncertainty. They chose the optimal decision mode by maximizing expected value.

In cases where ROA is intractable, decision rules provide a useful approximations of the optimal solution [205]. Generic decision rules also provide guidance for human decision-makers, since they are intuitive and straight-forward. Decision rules can take the form of decision trees [125], decisions networks [220], and logical conditions. In Decision Tree Analysis (DTA), decision-making depends on how the discretized, uncertain variables unfold over time. The conditions for making a particular decision depends on the particular tree limb. Decision rules can also take the form of logical “if, elseif, then” statements, where actions are triggered based on estimated state vectors [208]. These types of decision rules, which can be tuned with experimentation, closely resemble human decision-making as uncertainty scenarios are revealed over time. Lastly, decision networks characterize decision-making with chance and decision nodes in a time-expanded network [220]. This approach allows the user to consider the evolution of the architecture alongside the evolution of decision-making.

There are four generic types of decision rules: zero-order, condition-go, linear, and constant [205]. Conditional-go decisions with the “if, elseif, then” format are based on an estimate of future conditions and past information. While conditional-go rules don’t provide optimal solutions, they provide simple, useful, and practical guidelines for decision-making.

This thesis contains a series of conditional-go decision rules with parametric, tunable multipliers for an array of individual-satellite-level decisions, satellite-fleet-level decisions, and servicing-infrastructure-level decisions. The algorithms for all decision rules can be found in Appendix B and are explained in greater detail in the sections below.

Satellite-Level Decision Rules

The simulation implements comprehensive decision trees for each satellite experiencing failure, retirement, or collection events. This logic is encapsulated in the `DecisionTree` class, which evaluates all available options based on satellite state, capabilities, and technology obsolescence. Figure 3.13 illustrates the decision tree for individual satellite decision logic.




Figure 3.13: Satellite-Level Decision Tree Logic

The decision function directs each satellite through a decision process based on its location (in space or on Earth), operational status (retired or failed), obsolescence level, serviceability (refuelable or repairable), and available infrastructure (warehouse capabilities). For satellites in space, three main pathways are evaluated. If a satellite is retired, not obsolete, and refuelable—with refueling infrastructure available—the function compares the NPV of refueling versus launching a new satellite. If refueling is more valuable, the satellite is marked for in-space refueling; otherwise, it proceeds to end-of-life evaluation. A similar process applies for satellites that are obsolete but repairable, or failed but still repairable, provided repair infrastructure exists. If repair NPV exceeds launch NPV, the satellite is serviced; if not, it moves to the next step.

When in-space servicing is not possible or not economically viable, the satellite is evaluated for deorbit or Earth return. This function compares NPVs for Earth return and refurbishment against launching new assets. A tuned 1.1 multiplier biases toward Earth return when its value justifies it. If return is preferable and feasible, the satellite is scheduled for recovery; otherwise, it is deorbited and removed from service.

NPV Calculation Framework The standardized NPV calculation process is as follows:

1. **Determine time horizon:** Calculate years until obsolescence = T_{obs} – obsolescence age, where T_{obs} is determined by 1/obsolescence rate

2. **Project annual cash flows:** For each year until the satellite becomes obsolete:

$$\text{Cash Flow}_{\text{year}} = \text{Yearly Revenue} \quad (3.17)$$

3. **Apply periodic refueling costs:** Every design life years (typically 5 years), deduct refuel cost from that year's cash flow

4. Calculate Net Present Value:

$$NPV = \sum_{year=1}^{T_{obs}} \frac{\text{Cash Flow}_{year}}{(1 + r_d)^{year}} - \text{Initial Service Cost} \quad (3.18)$$

where r_d is the discount rate. This thesis assumes a constant, risk-free discount rate of 0.03.

5. **Add policy incentives:** $NPV_{final} = NPV + \text{Reuse Subsidy}$ (if applicable)
6. **Compare to the NPV of a new satellite:** Conduct similar calculations for a new satellite with the same time horizon, including relevant costs, such as launch and deorbit costs

Satellite Fleet-Level Decision Rules

RPO Upgrades The RPO upgrade decision rule determines when a customer should begin upgrading newly manufactured satellites to be capable of rendezvous and proximity operations to support self-delivery to the warehouses. This decision is periodically evaluated (every 3 months) based on the evolving economic trade-offs between building new satellites and refurbishing existing ones.

The upgrade rule activates under the following conditions: RPO decisions are enabled (flexible decision-making is enabled), RPO upgrades have not already been adopted, and in-orbit warehouse infrastructure exists. Once these preconditions are satisfied, the algorithm performs a cost-benefit analysis.

Two main cost comparisons are made: (1) the cost of building and launching new satellites versus (2) the total cost of servicing existing satellites, including refueling, repairing, refurbishment, Earth return, and ADR (Active Debris Removal) operations. These servicing costs are adjusted based on whether the customer has enabled refuel or repair upgrades (which apply cost streamlining effects for Earth-based service and enable space-based service) and whether the warehouse infrastructure has refueling

and repairing capability. The expected number of satellites eligible for refurbishment is estimated based on the remaining mission time, orbital configuration, and upgrade adoption percentages.

Market readiness and warehouse utilization are also factored in, along with policy incentives such as rebates or government subsidies. If the cost savings from servicing (relative to new satellite production) exceed the total RPO upgrade investment—computed as a fixed cost plus a per-satellite upgrade cost scaled by the expected satellite count—then the customer switches to producing RPO-capable satellites.

The percentage of satellites upgraded for Rendezvous and Proximity Operations (RPOD) changes dynamically based on system readiness and customer configuration. Before in-space refueling capability is enabled, the upgrade decision relies on a pre-refuel probability set by the user. This reflects a cautious approach, where only a small fraction of new satellites are upgraded for RPOD due to limited servicing support. Once refueling capability becomes available, the percentage increases to a post-refuel-upgrade value, reflecting increased confidence in the satellite servicing infrastructure.

This upgrade percentage directly influences how many satellites are considered eligible for RPO upgrades. The number is calculated as the product of remaining five-year operational cycles, the number of orbital planes, satellites per plane, and the factor value. As a result, the proportion of RPO-upgraded satellites evolves over time, influenced by the existing number of RPO satellites.

Function Start




Figure 3.14: RPO Upgrade Decision Rule Tree
154




Figure 3.15: Satellite Refuelability/Repairability Decision Tree
155

This flexible structure allows the RPO adoption rate to scale appropriately. When there are fewer than (number of warehouses)*(warehouse capacity) active satellites with RPO-capability, the percentage of new satellites having RPO-capability is 25%. Otherwise, the percentage of new satellites with RPO-capability is 15%.

The result is a dynamic, data-driven approach that activates RPO upgrades only when they are economically justified and logically supported by infrastructure and market conditions. Figure 3.14 illustrates the decision tree for the RPO upgrade decision.

Satellite Refuelability and Repairability Upgrades Customer decisions about investing in satellite capability upgrades (refuelability or repairability) employ Monte Carlo breakeven analysis embedded within the decision tree, illustrated in Figure 3.15

The upgrade decision process begins with the setup phase, which defines the fixed R&D cost, per-satellite upgrade cost, projected years remaining, and current collection rate. The process then conducts a Monte Carlo simulation running trials (10,000), where each trial samples the collection rate from a normal distribution with mean equal to the historical rate and standard deviation equal to 2 collected satellites. Negative rates are clipped to zero, and the total refurbishments over remaining years is calculated as rate multiplied by years.

Future collection rates are modeled as normally distributed, $R \sim \mathcal{N}(\mu_R, \sigma_R^2)$, to capture operational uncertainty in satellite collections per year. The normal distribution is justified by the Central Limit Theorem, as collection rates aggregate effects from multiple independent factors including market conditions, technical performance, and orbital dynamics. Using Monte Carlo simulation with $N = 10,000$




Figure 3.16: Add Warehouse Decision Rule Tree

samples, breakeven conditions are evaluated across the distribution of possible futures.

For the savings calculation, each trial computes the savings per satellite as the difference between the sum of new satellite cost and launch cost minus the refurbishment cost. Discounting is applied using the formula:

$$\text{Total Savings} = \sum_{t=1}^T \frac{\text{Savings} \times \text{rate}}{(1 + r_d)^t} \quad (3.19)$$

The investment calculation determines the total investment required through:

$$\text{Total Investment} = \text{Fixed R\&D} + (\text{Per-Sat Upgrade Cost} \times \text{Total Refurbishments}) \quad (3.20)$$

Breakeven determination identifies the minimum number of refurbishments where total savings are greater than or equal to total investment. Finally, the decision rule approves upgrade investment if the mean expected refurbishments exceed the breakeven threshold and the current collection rate exceeds the minimum threshold.

Servicing-Infrastructure-Level Decision Rules

Add Warehouse The decision rule that evaluates whether to add new orbital warehouse capacity is based on economic viability, comparing potential revenue from unserviced satellites against the cost of deploying new infrastructure.

The warehouse deployment decision process, illustrated in Figure 3.16, evaluates both potential revenue and associated costs to determine when adding a new warehouse is justified. Revenue is estimated from unserviced satellites, including failed satellites requiring Active Debris Removal (ADR), retired satellites that are refuelable or not, and additional income from penalty or rebate recovery depending on the active policy scheme. Costs include the base warehouse construction cost (which may be adjusted by multipliers depending on manufacturing cost trends and learning curve

progress), launch costs tied to payload requirements, and possible upgrade costs for refueling or repair capabilities. A new warehouse is added when projected revenue per existing warehouse exceeds the adjusted cost, when early deployment thresholds are met (such as having fewer than three warehouses), or when sufficient subsidization funds are available. The decision logic is also sensitive to policy configurations, which may influence costs through premiums, taxes, or fees, or offer financial subsidies to




Figure 3.17: Upgrade Warehouse Decision Rule Tree

prompt infrastructure investment through government intervention.

Upgrade Warehouse The warehouse upgrade decision process, depicted in Figure 3.17, operates on a monthly basis and determines whether to enhance all existing warehouses with capabilities for satellite refueling and/or repair. The decision begins by evaluating the need for a refueling upgrade. If there are no refuelable satellites or if the warehouses have already been upgraded for refueling, no action is taken. Otherwise, the system calculates a cost delta, defined as the difference between the cost of purchasing new satellites and the cost of servicing them, including expenses such as refurbishment, refueling, Active Debris Removal (ADR), and potential revenue from rebates or refunds (if applicable under the customer's policy). It then computes the fixed upgrade cost, which may be partially offset by government subsidies when subsidy mode is active and sufficient funds are available. If the cost delta exceeds the adjusted upgrade cost, scaled by a customer-defined multiplier, the system upgrades all warehouses to support refueling operations.

Once the refueling decision is resolved, the system proceeds to assess the need for a repair upgrade. This process mirrors the refueling logic: it checks for the presence of repairable satellites and whether warehouses have already been upgraded. If not, it recalculates the cost delta using updated service cost models that include repair expenses, and then compares this to the adjusted repair upgrade cost (again accounting for any subsidies). If the savings justify the investment, all warehouses are upgraded for repair capabilities.

Throughout the process, the logic incorporates policy-driven incentives. These include subsidy schemes that reduce upgrade costs and revenue enhancements from recovered penalties (for failed satellites) or policy-driven rebates tied to satellite operational lifetime. Together, these factors ensure that warehouse upgrades occur only when economically beneficial, strategically necessary, and supported by broader pol-

icy objectives.

3.5.1 Experimental Tuning Process

Tuning parametric decision values within the conditional-go decision trees follows the guidance provided by Lin et al. [208]. Rather than optimize decision rules, they took a trial and error approach to arrive at decision rules that provided sufficient downside tail reduction and upside tail extension in their Value-at-Risk/Gain (VARG) plots. This thesis takes a similar approach.

The parameter tuning process begins by identifying which multipliers or threshold values require improvement. Multiple simulation scenarios are designed and executed, systematically varying each parameter across a range. Each parameter setting is evaluated based on objective metrics such as total cost. Statistical validation is then applied using appropriate tests to determine which parameter settings produce statistically significant differences in outcomes and whether observed improvements represent reliable trends rather than random variation. Parameter values that reduce costs while meeting statistical significance thresholds are selected. These validated parameter values are then applied across all subsequent simulation iterations, ensuring consistent decision-making behavior throughout the analysis.

Ideally, the decision rules should be optimized in order to obtain the best possible results from the framework. However, the framework is intended to contain mid-fidelity models that simply screen and rank the potential of the proposed CONOPs and strategies. Optimizing these decision rules is recommended for future work, but is outside the scope of this thesis.

Having defined the decision rules for each of the satellite and level options, the following section explains how the decision rules are contained within the simulation and react to events and the evolving uncertain landscape.

3.6 Formulation Question 4: DES Component Modeling

Formulation Question 4

How to model the various components of the DES?

3.6.1 Formulation Question 4.1: Object Modeling

Formulation Question 4.1

How to model constellation and infrastructure objects?

The DES is structured using Python's SimPy library for discrete event management and implements an object-oriented architecture to represent the complex interactions between Active Debris Removal (ADR) vehicles, warehouses, customer satellites, and stakeholder entities. Over the simulation duration of 30 years, the framework explores a broad range of plausible scenarios through Monte Carlo simulation to enable comparisons of robustness and relative performance between different operational concepts and flexible options.

Object Class Architecture

The DES is structured using Python's SimPy library for discrete event management and implements an object-oriented architecture to represent the complex interactions between Active Debris Removal (ADR) vehicles, warehouses, customer satellites, and stakeholder entities. Over the simulation duration of 30 years, the framework explores a broad range of plausible scenarios through Monte Carlo simulation to enable comparisons of robustness and relative performance between different operational concepts and flexible options. The various object classes and their interactions while Figure 3.19, later in the section, provides the comprehensive logic flow diagram. The simulation implements six primary object classes that encapsulate the attributes of the system elements. Figure 3.18 provides a top-level methodology illustration.




Figure 3.18: Top-Level Methodology Diagram

Satellite Class

Each satellite object represents an individual spacecraft within the customer constellation. The baseline constellation resembles the OneWeb architecture with 18 orbital planes, 36 satellites per plane, at 1200 km altitude with 86.4-degree inclination. The satellites track their orbital state through position parameters. Physical properties include dry mass and fuel mass, while operational state management encompasses active, failed, retired, or deorbited status with associated event triggers. Each satellite maintains capability flags as boolean indicators for refuelable and repairable configurations, and tracks technology obsolescence through obsolescence age (years) and rate, which influence revenue contribution if the satellite's age exceeds the time to obsolescence in that scenario. Service pricing is dynamically calculated based on the satellite's willingness-to-pay for collection, determined through NPV analysis. Lifecycle event management includes start date, time to failure (exponentially distributed), design life, and retirement triggers, while assignment tracking monitors status for ADR collection assignment or other mission types.

The failure event mechanism operates by assigning each satellite a time to failure upon initialization, sampled from an exponential distribution with mean time between failures (MTBF) derived from the constellation-level failure rate. When simulation time reaches the assigned failure time, the satellite's status changes to 'failed' and a decision event is triggered to evaluate disposal options. Similarly, the retirement event mechanism triggers when a satellite completes its design life, changing status to 'retired' and initiating the decision-making process.

ADR Vehicle Class

ADR vehicles are responsible for collecting failed and retired satellites and transporting them to warehouses. The class implements collection logic with learning-based success probability. Each ADR object tracks its current plane location or warehouse location, maintains physical properties including dry mass, maximum payload

capacity (mass-constrained), fuel mass, and specific impulse for propulsion calculations. The operational status indicates whether the vehicle is available, unavailable, or in-transit. The ADR maintains a collected satellite inventory as a list of currently held satellites with total collected mass tracking. It also adjusts its collection success probability through a learning curve model.

The collection success mechanism follows a learning curve model representing improving operational proficiency, expressed as:

$$P_{success}(n) = 1 - (1 - S_0) \cdot n^{-\lambda} \quad (3.21)$$

where n is the cumulative number of satellites collected (starting from 1), S_0 is the initial success probability (scenario-specific random value), λ is the learning exponent controlling improvement rate (scenario-specific), and the probability plateaus after N_{max} collections. When an ADR attempts to collect satellites, each satellite undergoes an independent Bernoulli trial with success probability $P_{success}$. Successfully collected satellites are added to the ADR inventory and removed from the unassigned satellite pool, while failed collection attempts leave satellites available for future attempts.

The drift and service process operates through the primary process function `drift_and_service`, which monitors the ADR's current orbital plane location and evaluates collection opportunities when entering a new plane. The decision to initiate collection is governed by multiple constraints including failure priority, revenue potential, fuel feasibility, capacity constraints, and temporal alignment. Priority is placed on failed satellites; if there is a failed satellite within a 5 year drift-time window of the ADR's present location, the ADR vehicle won't collect any other satellites until it collects the failed satellite. Revenue potential requires that the sum of service prices for satellites in the plane must justify operation costs. Fuel feasibility demands that total mission fuel requirement must not exceed available fuel reserves. Capacity constraints ensure that total collected satellite mass cannot exceed ADR payload ca-

pacity (3 satellites). Temporal alignment requires that mission duration should align with warehouse drift periods to enable timely delivery.

The collection decision algorithm is implemented through the `ADR_cost_and_benefits` helper function, which simulates potential collection scenarios using deep copies of system objects to avoid modifying actual state during evaluation. The algorithm first filters satellites in the current plane by status (failed or retired), assignment (unassigned), and location (customer altitude), then sorts satellites by proximity to minimize in-plane maneuver costs. It iteratively adds satellites to a potential collection list while checking fuel constraints (cumulative fuel for rendezvous with each satellite plus return to warehouse), mass constraints (total collected mass including current satellite), revenue constraints (cumulative service prices minus operation costs), and time constraints (mission completion before next warehouse drift arrival). The algorithm calculates fuel requirements using Hohmann transfer equations for each satellite rendezvous and final warehouse return, compares total mission costs (fuel, operations) to the value of servicing the satellite to determine feasibility, and returns the selected satellite list if feasible or an empty list otherwise.

Maneuver execution is handled through the `maneuver` function, which executes the collection sequence. For each satellite in the selection list (in proximity order), the function calculates fuel required for rendezvous using the rocket equation, deducts fuel from ADR fuel mass, calculates maneuver time based on Hohmann transfer duration, advances simulation time by maneuver duration, and applies collection success probability via Bernoulli trial. If successful, the satellite status is updated to 'collected', added to ADR inventory, and ADR payload mass is incremented. If failed, the satellite is left unassigned and the failed collection attempt is logged. Both ADR and satellite orbital parameters are updated. After the final satellite, the function invokes `go_home` to return to the nearest available warehouse.

The warehouse return logic is implemented through the `go_home` function, which

identifies the warehouse that will next drift into the ADR’s current plane, calculates the drift time using J2 perturbation rates, performs a Hohmann transfer to parking orbit altitude, waits for RAAN alignment, and initiates the docking sequence with the warehouse.

Satellites with RPO (Rendezvous and Proximity Operations) capability can autonomously maneuver to the vicinity of the orbital warehouse but still require an ADR vehicle for final collection and last-mile delivery. This approach eliminates the need for expensive hybrid propulsion systems on each satellite, significantly reducing per-unit RPO upgrade costs while maintaining operational flexibility.

For these local ADR collection operations, the simulation makes three key assumptions. First, it assumes that the ADR vehicle can collect 3 RPO-capable satellites within a 7-day period. This assumption is justified because RPO-capable satellites can autonomously approach the warehouse and maintain station-keeping in its vicinity, allowing the ADR vehicle to efficiently service multiple spacecraft during a single sortie. The 7-day window accounts for coordination, approach sequencing, and safe capture operations for three satellites, reflecting realistic operational tempo without overtaxing the ADR vehicle. Second, the simulation assumes the ADR operational cost for RPO-capable satellites is 50% of the cost for non-cooperative satellite operations. This reduction is justified because RPO-capable satellites eliminate the most challenging and fuel-intensive phases of ADR: target approach, matching of tumbling or uncontrolled trajectories, and long-distance rendezvous. By arriving cooperatively at the warehouse vicinity, these satellites reduce ADR complexity to simple station-keeping and final capture maneuvers. Third, the simulation assumes that one local collection trip expends 20% of the ADR vehicle’s fuel capacity. This is justified because local collection operations involve minimal delta-V requirements compared to non-cooperative ADR missions, with the ADR vehicle only performing short-range maneuvers, attitude adjustments, and controlled returns to the warehouse rather than

extensive pursuit trajectories. While these assumptions are simplified for simulation purposes, they reasonably capture the operational advantages of cooperative RPO: reduced mission complexity, lower fuel requirements, and improved collection efficiency compared to traditional non-cooperative ADR operations.

Warehouse Class

Warehouses occupy a parking orbit altitude (lower than customer altitude) with the same inclination as the satellite constellation. They store spare satellites and receive collected satellites from ADR vehicles. Using the optimized capacity from Jakob et al.’s multi-echelon sparing paper, this framework assumes that warehouses can store 35 total satellites [192]. Each warehouse tracks its plane location by monitoring RAAN evolution due to J2 perturbation. Container management involves separate tracking of spare satellite count and collected, refurbished, and obsolete satellite count. Physical properties include dry mass, payload mass, fuel mass, and maximum capacity. Service capabilities are indicated by boolean flags for refueling equipment and repair equipment availability. ADR fleet management tracks which ADR vehicles are currently attached. Resupply event tracking monitors critical spare level and fuel level thresholds that trigger resupply requests, with requested warehouse upgrades delivered along with resupply missions.

The drift and replenish process is managed through the `drift_and_replenish` process function, which continuously propagates the warehouse’s RAAN. When the warehouse’s RAAN enters a new plane’s RAAN range (within tolerance), the process checks the current spare satellite count in that plane and compares it against the target spare level for the plane. If a deficit exists, the process initiates `release_new_sat` for each needed spare. If the warehouse contains refurbished satellites with satisfactory time to obsolescence and the plane has satellites retiring soon, the warehouse releases new active satellites along with spare satellites. Satellites are released with appropriate timing delays for orbital positioning, the warehouse spare inventory is

decremented, and the active satellite count in the plane is incremented.

The satellite reception mechanism is implemented when an ADR vehicle arrives at the warehouse through the `receive_satellite` function, which manages satellite transfer. The function first verifies that the warehouse has sufficient capacity (current collected plus incoming must be less than or equal to maximum capacity). If warehouse capacity is insufficient, the `monitor_collection` process waits for spare releases to create space, the ADR vehicle remains docked, and reception resumes.

The resupply process is managed through the `warehouseResupply` function, which monitors spare levels continuously. When onboard spare count drops below the critical threshold, the function triggers a resupply event, calculates required spare satellites, fuel needs, and any queued upgrades, invokes the launcher class to select the appropriate launch vehicle based on total payload mass, initiates a launch delay period (processing time plus stochastic component), and upon launch completion replenishes the warehouse with new spares, fuel, and upgrade modules as needed. If collected satellites are marked for Earth return, the function then initiates a return mission with associated transport costs using the same resupply launch vehicle.

Customer Class

The customer object represents the satellite constellation operator and manages financial accounting and facilitates strategic decisions when certain events trigger. Financial tracking is maintained through separate cost accumulators for different categories including launches, upgrades, operations, lost revenue, collisions, and servicing. Revenue tracking captures the time-varying revenue stream based on operational satellite count and technology obsolescence factors. The infrastructure upgrade state is monitored through boolean flags tracking satellite refuelability, repairability, and upgrade timing. Policy fund management tracks subsidy funds for policy intervention scenarios. Decision parameters store discount rate, service price multipliers, obsolescence thresholds, and upgrade costs. The class regularly updates its counts of

active, failed, retired, and spare satellites across all planes.

The decision process mechanism is implemented through `customerDecisionProcess`, a SimPy process that waits for satellite decision events or warehouse delivery events. When triggered, the process identifies which satellites have triggered events (status changed to failed/retired or delivered to warehouse), invokes the satellite-level decision tree for each event satellite, executes the chosen action (queue for service, assign for collection, or deorbit), updates financial accounting based on decision outcome, and resets the decision event for the next trigger.

Launcher Class

The launcher manages all launch operations and vehicle selection. It maintains a launch vehicle dictionary storing performance, capacity, and economic parameters for seven launch vehicles, tracks scenario-specific time-varying cost trajectory, manages launch timing including processing time, stochastic delay components, and stores the uncertain multiplier for return that determines how expensive Earth return missions are compared to the regular launch cost.

The vehicle selection mechanism is implemented through the `best_rocket_for_payload` function, which performs cost-optimal vehicle selection. Given the required payload mass and current simulation time as inputs, the function filters available vehicles where payload capacity is greater than or equal to required mass and operational maturity time is less than or equal to current time. If no vehicles satisfy constraints, the function returns a failure indicator. From valid vehicles, it selects the one with minimum payload capacity (smallest vehicle that meets requirements) and returns the vehicle name, maximum capacity, second stage mass, dry mass, and baseline cost per kilogram. The actual mission cost is calculated based on the payload and the given vehicle's estimated launch cost (\$/kg) at the start of the simulation multiplied by the change in launch cost since then. This framework assumes that remaining payload capacity in the selected launch vehicle, should it exist, is sold to another

entity. The cost of this empty payload space if not accounted for.

Launch execution occurs when a resupply or deployment is needed. The process calculates total payload (satellites plus fuel plus upgrades plus ADRs), invokes the vehicle selection algorithm, samples launch delay from an exponential distribution, calculates total launch time as processing plus delay, yields simulation for launch time duration, and upon completion delivers payload to destination, updates emissions tracking, and updates cost accounting.

Infrastructure Class

The Infrastructure class coordinates adaptive, fleet-level decision-making across the Collection-as-a-Service (CAAS) system lifecycle, managing infrastructure investments in response to evolving operational conditions and policy incentives. Core responsibilities include warehouse and satellite fleet management, which coordinates warehouse upgrade decisions and the deployment of new warehouses, as well as satellite fleet modernization, which oversees constellation-wide decisions to upgrade satellites with refuelable/repairable/RPO capabilities.

The Infrastructure class implements concurrent decision-making processes that evaluate infrastructure needs at different times in the simulation. Warehouse upgrade evaluation occurs at monthly intervals and begins by identifying satellites requiring service based on fleet composition, specifically refuelable or repairable satellites currently collected in warehouses or active in orbit. The process calculates aggregate service revenue potential across the fleet, accounting for a practical number the warehouse could service, and applies decision rules: if total refuelable satellite cost savings are greater than or equal to refuel upgrade cost, a refuel capability upgrade is queued; if total repairable satellite cost savings are greater than or equal to repair upgrade cost, a repair capability upgrade is queued. Queued upgrades are delivered during the next warehouse resupply mission.

The satellite fleet upgrade decision undergoes periodic Monte Carlo evaluation.

The process tracks historical collection/refurbishment rates across the entire constellation, projects future service demand over the remaining mission lifetime (typically 20 years), and runs 10,000 Monte Carlo trials with discounted cash flow analysis to compare the cost of upgrading new satellites to be refuelable/repairable (including R&D and per-satellite multipliers) against savings from refurbishment versus replacement (accounting for policy rebates/subsidies). The decision is triggered if the projected collection rate exceeds the break-even threshold. The upgrade percentage scales with collection rate, with a minimum percentage of satellites upgraded set as a function of warehouse number, and incorporates risk preferences through a parametric multiplier.

New warehouse addition undergoes periodic evaluation with a time constraint. The process aggregates uncollected satellite service prices (revenue in waiting) and calculates cost including base warehouse cost, launch costs, learning curve effects (cost multiplier decreases with each warehouse), and any required capability upgrades (refuel/repair). The decision rule adds a warehouse if cost-savings-per-warehouse is greater than or equal to total infrastructure cost. A time constraint prevents new warehouses from being added after year 20, allowing sufficient operational time for cost recovery. The process prevents concurrent warehouse additions through the `whADDinprogress` flag. The function also prevents more warehouses than there are orbital planes, setting a practical limit on warhouse infrastructure expansion.

Continuous monitoring processes include service price updates, which recalculate satellite-specific service prices when satellites experience milestones and collision cost calculation, which tracks exponentially increasing collision risk costs based on orbital congestion that is assumed to increase with time.

The Infrastructure class accounts for fleet-level decision-making by balancing reactive responses (based on current fleet state) with predictive analysis (Monte Carlo projections), incorporating policy incentives directly into cost calculations (subsidies reduce upgrade costs, rebates increase refurbishment value), using learning curves to

model decreasing infrastructure costs with scale, coordinating timing constraints to ensure infrastructure investments have sufficient operational lifetime for positive ROI, and synchronizing satellite-fleet and warehouse-fleet perspectives through shared service price signals. This multi-scale decision framework enables the simulation to capture emergent behaviors where infrastructure evolves adaptively in response to uncertain satellite failures, changing technology costs, and policy interventions—mimicking realistic stakeholder decision-making in nascent space servicing markets.

Emissions Class

The Emissions class tracks atmospheric pollution from space operations by monitoring reentering mass that contribute to NO_x emissions. The class distinguishes between two primary reentry pathways that have different environmental impacts. Core properties include `massDown`, representing reusable mass that reenters intact via controlled descent, such as reusable second stages (measured in kg); `deorbitedMass`, representing mass that burns up during atmospheric reentry, including satellites and debris (measured in kg); and `nonreusableSecondStageMass`, representing mass of expendable rocket stages that deorbit after use and burn up in the atmosphere (measured in kg). These properties are updated every time a satellite deorbits or a launch vehicle reenters the atmosphere.

SimPy Event Management and Process Orchestration

SimPy orchestrates the discrete event simulation through event-driven processes, illustrated in the comprehensive flow diagram in Figure 3.19. The simulation manages multiple concurrent process functions that yield control to the SimPy environment and wait for specific events or timeouts. Primary event types include satellite failure events, which are pre-scheduled at initialization based on exponentially distributed failure times and trigger status change and decision tree invocation; satellite retirement events, which are triggered after satellite completes design life and initiate end-of-life decision process; warehouse resupply triggers, which are activated when spare

count drops below critical threshold and initiate launch request; warehouse delivery events, which are triggered when an ADR vehicle docks with a warehouse and initiates satellite transfer; customer decision events, which are AnyOf events that trigger when any satellite has a decision to make or when warehouse deliveries occur; infrastructure evaluation events, which provide periodic assessments (typically monthly) for upgrade and expansion decisions; and parameter update events, which provide scenario-specific updates to uncertain variables at predetermined intervals.

The simulation implements standard SimPy process function patterns. The infinite loop with event waiting pattern is used for warehouse drift, infrastructure evaluation, and parameter updates. The AnyOf event pattern allows the customer decision process to respond to multiple potential event sources efficiently. The sequential process with yields pattern is used for ADR collection missions and launch operations.

Throughout execution, the simulation continuously records time-stamped financial transactions, operational event logs (collections, services, launches), system state snapshots at key intervals, decision outcomes and associated NPVs, infrastructure deployment events and timing, policy intervention applications (subsidies, refunds, fines), and sustainability metrics (number of satellite collections/refurbishments, number of deorbits, and cumulative deorbited mass).

Data is collected in memory during simulation execution and written to structured output files (Excel or CSV) at completion for post-processing analysis. Figure 3.19 shows the entire logical flow diagram for the Discrete Event Simulation.




Figure 3.19: Methodology Logical Flow Diagram

3.6.2 Formulation Question 4.2: Cost Modeling

Formulation Question 4.2

How to model costs?

This work combines cost models and estimates from several sources. A summary of all simulation costs is included in Table 3.4. Specific cost assumptions and their justifications are available in Table 3.5.

Table 3.4: Summary of Simulation Costs

Cost Type	Details
Initial cost	Includes the cost of existing satellites, warehouses, and ADR vehicles at the start of the simulation, including any applicable upgrades.
Cost to directly replace failed satellites	Triggered when no in-plane spare or nearby warehouse is available. Includes cost of the satellite, its fuel, and a single-satellite launch.
Cost to launch ADR mission directly from Earth	Applied when no ADR vehicle is available nearby. Includes ADR hardware, launch cost, and operation cost.
Cost to replace planes of satellites	Scheduled batch replacement of aging planes. Includes costs of satellites, fuel, and launch for all replacements.
Cost to resupply warehouses	Includes spare satellites, satellite and warehouse fuel, new ADRs (if ordered), upgrades, and return cost for old satellites if sent to Earth.

Cost Type	Details
Cost to service satellites on Earth	Incurred when satellites return and are refurbished or repaired based on NPV evaluation. Includes servicing and recommissioning costs.
Cost of lost revenue due to obsolescence or lost coverage	Revenue loss due to satellites being down or obsolete. Obsolescence reduces utility; gaps in coverage incur cost based on revenue loss per day.
Cost associated with collision avoidance or events	Incurred when collision avoidance maneuvers or actual collisions occur, as defined by an exponential cost function that scales with the scenario's temporal progression (reflecting increasing orbital congestion), the population of non-maneuverable satellites, and their orbital residence time
Cost of ADR operations in space	Cost to operate ADR missions within space (e.g., satellite retrieval and towing to warehouse).
Cost to upgrade satellites	Applied when customers invest in satellite upgrades, such as refuelability or repairability.
Cost to upgrade warehouses	Applied when a warehouse is upgraded (e.g., to include fuel or repair capabilities).
Cost to service satellites in space	Incurred when satellites are repaired or refueled on-orbit using warehouse capabilities.
Cost to deorbit satellites	Applied when a satellite is actively deorbited instead of refurbished or reused.
Cost of adding new warehouses	Applied when customers expand capacity by deploying new orbital warehouse facilities.

Table 3.5: Initial Cost Assumptions (Year: 2025)

Variable Name	Value	Justification
Satellite Unit Cost	\$900,000	Based on estimated OneWeb satellite cost [191]
Earth Refurbishment (General)	\$200,000	Includes labor, testing, parts and upgrades, and facilities. It's assumed that refurbishing is roughly 20-40% of manufacturing a new satellite, based on the economics of refurbishing reusable rockets, which is 65% cheaper than launching new [221]
In-space/Earth Repair (Repairable Satellite Failure)	\$450,000	More effort than refurbishment, roughly half the cost of a new satellite
In-space/Earth Repair (Obsolescence, Repairable)	\$550,000	More expensive than repair due to payload upgrade and related testing
Earth Repair (Obsolescence, Non-repairable)	\$855,000	Requires more extensive labor to swap a payload in a satellite not meant to be repaired, roughly the same cost as a new satellite

Continued on next page

Table 3.5 – continued from previous page

Variable Name	Value	Justification
Satellite Refuel Upgrade, R&D	\$5,000,000	Redesign for mechanical and fluid interface, propulsion system adjustments, electrical & software integration, testing and certification
Satellite Repair Upgrade, R&D	\$9,000,000	Design modularization, standardized interfaces, on-board failure monitoring, testing, qualification
Warehouse Refuel Upgrade	\$8,000,000	Xenon storage and pressure management, refueling system with specialized docking, precision control systems, diagnostic monitoring, coordination between depot and satellite, training, certification, and testing
Warehouse Maintenance	\$15,000,000	Diagnostic systems and related sensors, robotic capability and tooling, spare parts, power systems for satellite battery test & recharge, specialized software, training, and testing

Table 3.5 – continued from previous page

Variable Name	Value	Justification
ADR Vehicle Cost	\$48,000,000	Based on \$16 million (USD) Astroscale contract that covers about 1/3 of vehicle cost [101]
ADR Operation	\$250,000	Includes the cost for ground support, GNC, RPOD, system checks, mission operations
ADR Operation for RPO satellite	\$83,250	Includes the cost for ground support, GNC, RPOD, system checks, mission operations. Less than classic ADR operations because RPO satellites are upgraded to be cooperative and maneuver duration is shorter
In-Space Refueling Operation	\$250,000	Includes the cost for ground support, controls, robotic operation, system checks, depreciation of hardware, labor, mission operations

Continued on next page

Table 3.5 – continued from previous page

Variable Name	Value	Justification
Warehouse Cost	\$100,000,000	Based on analogous servicing spacecraft budgets, such as DARPA RSGS [222] and MEV-2 [223], which have more capability than initial spare warehouse configuration. Can also be compared to the scale of large GEO satellites, such as Intelsat 10-02 [223], factoring out launch cost to GEO
Deorbit Cost (From constellation altitude)	initially \$100,000; set to 1/9 of spacecraft cost and varies proportionally	Based on NASA estimate of additional cost for medium satellite's extra propellant to immediately deorbit from 800 km ranges (between \$85,000 and \$425,000 [224])
Satellite to Warehouse RPO Cost	1/2 of present deorbit cost	conservative estimate based on estimated deorbit cost [224]
Deorbit Cost (From parking orbit altitude)	1/2 of present deorbit cost	conservative estimate based on estimated deorbit cost [224]

3.6.3 Formulation Question 4.3: Trajectory Modeling

Formulation Question 4.3

How to model trajectory?

High Thrust Maneuvering

This work assumes that all orbits are circular and all ADR orbital maneuvering can be represented with Hohmann transfers. The V_{Hohmann} for an in-plane maneuver is provided in the following equations [225]:

$$a_{\text{transfer}} = \frac{r_0 + r_1}{2} \quad (3.22)$$

$$v_0 = \sqrt{\frac{\mu}{r_0}} \quad (3.23)$$

$$v_{\text{transfer},0} = \sqrt{\mu \left(\frac{2}{r_0} - \frac{1}{a_{\text{transfer}}} \right)} \quad (3.24)$$

$$v_1 = \sqrt{\frac{\mu}{r_1}} \quad (3.25)$$

$$v_{\text{transfer},1} = \sqrt{\mu \left(\frac{2}{r_1} - \frac{1}{a_{\text{transfer}}} \right)} \quad (3.26)$$

$$\Delta v_{\text{Hohmann}} = |v_{\text{transfer},0} - v_0| + |v_1 - v_{\text{transfer},1}| \quad (3.27)$$

In these Hohmann transfer equations, the variables are: a_{transfer} for the semi-major axis of the transfer orbit, r_0 for the initial orbital radius, r_1 for the final orbital radius, v_0 for the velocity in the initial circular orbit, $v_{\text{transfer},0}$ for the velocity at periapsis of the transfer orbit, v_1 for the velocity in the final circular orbit, $v_{\text{transfer},1}$ for the

velocity at apoapsis of the transfer orbit, μ for the gravitational parameter (GM), and $\Delta v_{\text{Hohmann}}$ for the total delta-v required for the Hohmann transfer. Accounting for the ΔV of each maneuver, the necessary mass can be determined using the rocket equation.

$$\text{mass}_{\text{wet}} = \text{mass}_{\text{dry}} e^{\frac{\Delta V}{gT_{sp}}} \quad (3.28)$$

Low Thrust Maneuvering

This simulation incorporates collaborative maneuvering capabilities within its CONOPs, where satellites equipped with electric propulsion systems can conduct RPO with orbital warehouses while specialized ADR vehicles collect them for subsequent, last-mile rendezvous and docking operations.

Given the computational complexity of precisely modeling low-thrust orbital transfers, this framework employs a conservative analytical approach. Edelbaum's analytical approximation [226] provides a method for computing delta-V for combined altitude and plane changes between circular LEO orbits:

$$\Delta v = \sqrt{v_0^2 + v_1^2 - 2v_0v_1 \cos\left(\frac{\pi}{2}\theta\right)} \quad (3.29)$$

where v_0 and v_1 are the circular orbital velocities at the initial and target altitudes ($v = \sqrt{\mu/r}$), and θ represents the plane change angle.

However, to ensure satellites retain sufficient propellant reserves for deorbit operations in the event that warehouses are at capacity or RPO operations are unsuccessful, this framework adopts a conservative operational constraint: satellites do not actively change RAAN through propulsive maneuvers. Instead, warehouses are positioned at the same inclination as the origin constellation, and satellites rely exclusively on natural J2 precession to drift to the required RAAN. This eliminates the propulsive plane change component ($\theta = 0$), reducing the delta-V to pure altitude change:

$$\Delta v_{\text{total}} = |v_1 - v_0| \quad (3.30)$$

Under this constraint, the total transfer time is determined by RAAN drift time from natural J2 precession for circular orbits [225]:

$$\begin{aligned} n &= \sqrt{\frac{\mu}{a^3}} \\ \dot{\Omega} &= -\frac{3}{2} \frac{n J_2 R_E^2}{a^2} \cos i \\ t_{\text{RAAN}} &= \frac{|\Delta\Omega|}{|\dot{\Omega}|} \end{aligned} \quad (3.31)$$

The framework uses $t_{\text{transfer}} = t_{\text{RAAN}}$ to capture the reality that satellites must wait for favorable orbital geometry. This conservative approach ensures that warehouse transfers remain propellant-efficient. Maneuvers exceeding 5 years are deemed impractical, and satellites that would take that long are directed to deorbit instead.

This implementation reflects a deliberate design choice: by limiting propulsive maneuvers to altitude changes only and leveraging natural orbital mechanics for RAAN phasing. This ensures satellites maintain adequate fuel reserves to independently deorbit if they arrive at a warehouse that is at capacity or if RPOD operations fail, preserving responsible end-of-life disposal capabilities as a fallback option.

3.6.4 Formulation Question 4.4: Emissions Modeling

Formulation Question 4.4

How to model emissions?

There are a number of emissions associated with launching and deorbiting satellites. For the purposes of this thesis, however, we will focus on NO_x , since it has a sizable impact on ozone compared to other emissions [36] and has a rule-of-thumb emission estimate that is a function of spacecraft mass and reentry condition. The framework will calculate the kgs of NO_x produced for each infrastructure simulation

so the user can compare the effectiveness of CAAS in reducing NO_x emission. This will include NO_x produced by the rocket launches necessary to resupply the customer constellations and CAAS infrastructure elements.

In 1976, Park et al. approximated reentering spacecraft as flat plates and determined that they produce about 4.5-9% of its mass in nitric oxide. This approximation assumes that air flow in the shock layer is in chemical equilibrium. This rule of thumb was improved in the paper, “Equivalent-cone calculation of nitric oxide production rate during space shuttle reentry” [38]. These updated calculations remove the chemical equilibrium assumption and use plan area to approximate the NO_x mass flow rate produced from a returning space shuttle, finding NO_x production to be 2.6 times greater than the Park estimation. In their environmental assessment, Ryan et al. estimate that returning reusable components produce NO_x emissions equivalent to 17.5% of its mass [36]. For spacecraft that completely vaporize, including reentering debris and rocket stages that fall from 50km or higher, they estimate that it contributes 100% of its mass as NO_x emissions [36]. These estimates are consistent with the Larson et al. study [37] and NASA Space Shuttle reentry estimates [38].

The proportional relationship for total NO_x emissions used in this framework is:

$$NOx \text{ emissions} \propto 0.175 \cdot m_{\text{Returned to Earth}} + m_{\text{burned in the atmosphere}} \quad (3.32)$$

To estimate the mass of returning launch vehicles, the emission model uses the Tsiolkovsky rocket equation for stage separation and second-stage performance and uses publicly available information about existing or developing launch vehicles [227],[228],[229], [230],[231]. Where information is not publicly available, the model uses approximations, available in Appendix A in Table A.2. The second-stage velocity increment is:

$$\Delta V_2 = I_{sp_2} \cdot g_0 \cdot \ln \left(\frac{M_{sep}}{M_{2ndstage_{dry}} + M_{payload}} \right) \quad (3.33)$$

The separation mass (M_{sep}) is estimated as:

$$M_{sep} = M_{\text{launch}} \cdot \exp \left(-\frac{V_{sep}}{I_{sp_1} \cdot g_0} \right) \quad (3.34)$$

And the dry mass of the second stage ($M_{2ndstage_{dry}}$), which contributes to reentry emissions, is:

$$M_{2ndstage_{dry}} = M_{sep} \cdot \exp \left(-\frac{\Delta V_2}{I_{sp_2} \cdot g_0} \right) - M_{payload} \quad (3.35)$$

3.6.5 Formulation Question 4.5: Policy Modeling

Formulation Question 4.5

How to model and incorporate the effect of various policy schemes?

While numerous policies could influence the OOS market in LEO, this research examines eight distinct policy approaches, including five single-mechanism policies and three hybrid combinations. The selection draws from established environmental economics theory and builds upon policy recommendations specifically developed for space sustainability challenges. Each scheme addresses different market mechanisms while targeting the common goal of promoting economically viable on-orbit servicing infrastructure that can reduce atmospheric emissions from satellite operations.

The policy selection process recognizes that government intervention fundamentally impacts the costs and benefits of options available to constellation operators. As noted in space sustainability research, effective policy should help “tip the scales” on viable business cases where needed, supporting improvements in constellation sustainability while limiting penalties on satellite operators and costs to taxpayers. The selected schemes represent distinct approaches to internalizing environmental costs

and creating economic incentives for sustainable space operations.

Policy Scheme 1: OUF/Refund

Orbital Use Fees form the foundation for two policy schemes, reflecting their prominence in space sustainability literature. Rao and colleagues demonstrated that optimal orbital use fees could significantly increase space industry value while effectively mitigating collision risks, since satellite operators do not fully account for collision costs they impose on one another. Their recommended fee of \$14,900 per satellite, escalating to \$235,000 by 2040, provides empirical grounding for the OUF parameters tested in this research [153].

The first OUF policy scheme implements the deposit/refund concept proposed by Macauley [152]. This approach directly incentivizes satellite collection and refurbishment by creating immediate financial returns for responsible end-of-life management. The accumulated fee structure ensures that longer operational lifetimes generate larger refund values, encouraging both extended satellite life and eventual responsible disposal.

Policy Scheme 2: OUF/Subsidy

This policy follows the Roy and colleagues approach of using orbital use fee proceeds to subsidize on-orbit recycling infrastructure [154]. This scheme addresses the classic collective action problem where individual operators lack sufficient incentive to invest in infrastructure that benefits the environment. By pooling OUF revenues to subsidize warehouse infrastructure or satellite upgrades, this approach can overcome initial capital barriers that prevent economically efficient servicing infrastructure development.

Policy Scheme 3: Contingent Fine/Subsidy Framework

Policy Scheme 3 targets failed satellites specifically. The fine structure creates direct accountability for satellite failures while the refund mechanism maintains incentives for responsible collection. This dual approach ensures that operators who contribute to the pollution problem bear the costs while those who participate in remediation receive financial compensation. Proceeds from the fines form a subsidy fund that can provide rebates for refurbishments and/or subsidize elements of the infrastructure or satellite upgrades.

The contingent fine approach addresses concerns raised about purely ex-ante policies that rely on complex forecasting methods. By implementing penalties based on actual satellite failures rather than predicted risks, this scheme avoids the conservative assumptions and methodological challenges that industry stakeholders like SpaceX have criticized in proposed aggregate collision probability metrics.

Policy Scheme 4: Insurance Market Integration

Mandatory insurance requirements build upon established space law frameworks while addressing current insurance market failures. Research from the Aerospace Corporation suggests that on-orbit servicing and orbital insurance could develop a “unique, symbiotic relationship” since servicing availability would reduce both the number and severity of insurance claims [155].

This policy scheme addresses multiple market failures simultaneously. The mandatory insurance requirement creates universal coverage that internalizes satellite failure risks, while the premium elimination for serviced satellites provides direct economic incentives for operators to utilize available servicing infrastructure. Insurance companies would naturally become advocates for servicing capabilities and standardization, creating market-driven pressure for sustainable operations.

Policy Scheme 5: Progressive Revenue Generation

Policy Scheme 5 (Subsidy/Taxes) implements a progressive taxation approach that scales with constellation operator success. This time-progressive structure acknowledges that successful operators have greater capacity to contribute to industry-wide sustainability infrastructure while ensuring that rapidly growing companies maintain responsibility for their increasing environmental footprint. The approach follows Adilov and colleagues' recommendation for Pigouvian taxation that funds debris mitigation, adapted specifically for atmospheric emissions concerns [151].

The profit-based taxation mechanism avoids penalizing operators during vulnerable phases while ensuring that established, profitable operations contribute proportionally to sustainability infrastructure. This approach addresses concerns about international competitiveness and “forum shopping” by scaling obligations with financial success rather than imposing uniform burdens regardless of operator circumstances. Similar to other policy schemes, the taxes form the subsidy fund.

Policy Scheme 6: Direct Tax + Fine + Subsidy Hybrid

Policy Scheme 6 implements the progressive taxation of policy scheme 5 along with infrastructure subsidies and the fine mechanism of policy scheme 3, creating multiple avenues to build the subsidy fund and encourage sustainable operations. The dual funding mechanism provides both immediate fine-based revenues and growth-scaled tax contributions, creating robust financial support for servicing infrastructure while distributing costs across operational penalties and profitable operations.

Policy Scheme 7: OUF + Fine + Subsidy Hybrid

Policy Scheme 7 combines the OUF-based revenue generation of Policy 2 with infrastructure subsidies and the fine mechanism of policy scheme 3. Similar to policy scheme 6, this hybrid scheme leverages multiple revenue streams to maximize subsidy

funds available for infrastructure development. This approach tests whether combining complementary policy mechanisms yields synergistic benefits or creates excessive regulatory burden.

Policy Scheme 8: Insurance/Subsidy Integration

Policy Scheme 8 mimics the mechanics of insurance premiums, but orchestrated through a government fund that can provide infrastructure subsidies, creating a hybrid between risk-based pricing (Policy 4) and subsidy fund development. Unlike Policy 4 where premium waivers provide the only incentive for servicing, Policy 8 uses the proceeds to directly subsidize warehouse capabilities and satellite upgrades and reward refurbishments.

Government Class Implementation

The `theGovernmentClass` orchestrates policy implementation through conditional activation of process functions based on the selected policy scheme. This class tracks government financial flows and triggers policy mechanisms throughout the simulation lifecycle. Note that at the end of the 30-year simulations, for all policies but policy 4, the government subsidy fund is subtracted from the constellation operator's total cost in order to compare the policy's impact on total system costs.

Core Properties:

- **funds:** Government-held subsidy pool available for infrastructure support
- **annualOUF:** Per-satellite orbital use fee (Policies 1, 2, 7)
- **failurePenalty:** Fine levied on failed satellites (Policy 3, 6, 7)
- **annual_premium:** Insurance premium percentage of satellite value (Policies 4, 8)

- **tax_rate:** tax percentage parameter (Policies 5, 6)
- **refund_condition:** Determines refund trigger (1 = collection, 2 = refurbishment) (Policy 1)
- **rebate_for_refurbishment:** Provides a \$250,000 rebate for every refurbished satellite (Policies 2, 3, 5, 6, 7, 8)
- **initial_subsidy:** Subsidizes the initial cost delta between the CAAS configuration and the baseline (Policies 2, 3, 5, 6, 7, 8)

Process Function Activation Logic:

The Government class conditionally activates specific process functions based on the policy scheme:

- **Policy 1:** `OUF_collect() + refund()`
- **Policy 2:** `OUF_collect() + subsidize()`
- **Policy 3:** `subsidize()` (fine collection handled in customer failure events)
- **Policy 4:** `apply_premium()`
- **Policy 5:** `tax() + subsidize()`
- **Policy 6:** `tax() + subsidize()` (fine collection handled in customer failure events)
- **Policy 7:** `OUF_collect() + subsidize()` (fine collection handled in customer failure events)
- **Policy 8:** `apply_premium() + subsidize()`

Key Process Functions:

OUF_collect(): Executes annually, identifying active satellites that have not been refurbished (neither on Earth nor in space). Adds the annual OUF charge to customer costs and credits the same amount to the subsidy fund.

refund(): Event-driven process that triggers when satellites are collected (refund condition = 1) or refurbished (refund condition = 2). Refunds the accumulated OUF for the satellite's operational age by reducing both customer costs and government funds.

tax(): Executes annually with time-progressive tax rates calculated as:

$$\text{tax_percent} = \text{maxTax} \times \left(\frac{t}{30}\right)^{\text{tax_shape_parameter}} \quad (3.36)$$

where t is mission time in years and $\text{tax_shape_parameter}=1$. The tax is applied to operator profit and added to subsidy funds.

apply_premium(): Identifies active satellites without refurbishment history and charges annual premiums as a percentage of original satellite cost. Premium revenues are not pooled for subsidies in Policy 4 (private insurance) but are in Policy 8 (government-managed subsidy fund).

subsidize(): Event-driven process triggered by warehouse additions, upgrades, or satellite capability enhancements that enables infrastructure cost reductions if subsidy funds are available.

This architectural separation between policy mechanisms enables clean comparison of single-instrument versus hybrid approaches while maintaining computational efficiency through event-driven process activation.

3.6.6 Addressing Regulatory Gaps

These policy schemes directly address the regulatory gap identified by the Government Accountability Office, which noted that insufficient research on atmospheric emissions

prevents current policy development. While orbital congestion has received substantial policy attention, atmospheric pollution from space operations remains largely unregulated, falling outside both National Environmental Policy Act and Montreal Protocol jurisdiction. The selected schemes recognize that atmospheric pollution mitigation and orbital congestion management share fundamental characteristics, allowing proven orbital policy concepts to inform atmospheric emissions strategies. Each scheme provides mechanisms to reduce mass flux through the atmosphere by incentivizing satellite life extension, refurbishment, and responsible collection, therefore addressing atmospheric emissions uncertainty through precautionary approaches that provide economic benefits regardless of the ultimate impact of emissions on the upper atmosphere.

The diversity of approaches tested, from market-based fees to insurance integration to progressive taxation, ensures comprehensive evaluation of different intervention philosophies while maintaining focus on economically viable solutions that support rather than hinder space industry growth. This selection enables identification of scenario-dependent optimal policies that can establish sustainable servicing infrastructures under varying economic and technological conditions.

3.6.7 Policy Impact Integration in Cost-Benefit Analysis

The discrete event simulation incorporates policy interventions through a comprehensive government class that systematically tracks financial flows and modifies constellation operator decision logic. The simulation employs a conditional execution structure where policy scheme selection determines which government processes are activated.

Direct Cost Impacts

The simulation tracks policy-induced costs through multiple channels that modify the baseline constellation operator cost structure. Annual orbital use fees are calculated as:

$$\text{Annual OUF Cost} = N_{\text{active}} \times \text{OUF}_{\text{annual}} \quad (3.37)$$

where N_{active} represents the number of active satellites neither Earth nor space refurbished, and $\text{OUF}_{\text{annual}}$ is the policy-specified fee per satellite.

For mandatory insurance schemes, premiums are computed as:

$$\text{Insurance Premium} = \sum_{i=1}^{N_{\text{active}}} C_{\text{original},i} \times r_{\text{premium}} \quad (3.38)$$

where $C_{\text{original},i}$ is the original cost of satellite i and r_{premium} is the annual premium rate.

Progressive taxation follows a time-dependent structure:

$$\text{Tax Rate}(t) = r_{\text{max}} \times \left(\frac{t}{30}\right)^{\alpha} \quad (3.39)$$

where r_{max} is the maximum tax percentage, t is time in years, and α is the tax shape parameter controlling the rate of increase.

Benefit Tracking and Rebates

The simulation incorporates policy benefits through rebate mechanisms tied to specific satellite lifecycle events. For OUF refund schemes, rebates are calculated based on satellite operational duration:

$$\text{OUF Rebate} = t_{\text{age}} \times \text{OUF}_{\text{annual}} \quad (3.40)$$

where t_{age} is the satellite age in years at the time of collection or refurbishment.

3.6.8 Decision Logic Integration

Satellite Lifecycle Decisions

Policy impacts are embedded within the Net Present Value (NPV) calculations that drive satellite servicing decisions. The simulation modifies the baseline NPV to include policy-specific cash flows.

For new satellite procurement, the total lifecycle cost incorporates policy expenses:

$$C_{\text{total}} = C_{\text{satellite}} + (\text{OUF}_{\text{annual}} \times L_{\text{design}}) + (C_{\text{satellite}} \times r_{\text{premium}} \times L_{\text{design}}) \quad (3.41)$$

where L_{design} is the satellite design life.

Infrastructure Investment Decisions

Warehouse addition and upgrade decisions incorporate policy-generated subsidy funds. The simulation evaluates infrastructure investments by comparing revenue potential against costs modified by available subsidies:

$$\text{Net Infrastructure Cost} = C_{\text{infrastructure}} - \min(F_{\text{subsidy}}, C_{\text{infrastructure}}) \quad (3.42)$$

where F_{subsidy} represents the available government subsidy fund.

3.6.9 Fund Flow Dynamics

The government class maintains a dynamic subsidy fund (F_{subsidy}) that accumulates revenue from policy mechanisms and disburses funds for infrastructure support. Fund accumulation follows policy-specific rules:

$$\frac{dF_{\text{subsidy}}}{dt} = R_{\text{OUF}} + R_{\text{tax}} + R_{\text{fines}} + R_{\text{premiums}} - D_{\text{subsidies}} - D_{\text{rebates}} \quad (3.43)$$

where R terms represent revenue inflows and D terms represent disbursements.

3.7 Framework Integration and Experimentation

This chapter develops the necessary elements for an OOS flexibility framework that screens flexible options, addressing gaps from preexisting OOS flexibility frameworks, modeling the novel CAAS CONOPs and its elements, and leveraging concepts from flexible infrastructure research. Following a series of formulation questions that address the requirements for experimentation on hypotheses 1 through 3, this chapter presents the logic for developing a flexibility framework that leverages a discrete event simulation that includes decision rules for both the satellite constellation and the servicing infrastructure. These decisions address experimental requirements 1-3, presented in Figure 3.21. The following figure, Figure 3.22, illustrates the steps from concept formation for result, indicating that tuning is necessary to improve decision rule performance. Figure 3.20 illustrates how to use the framework to conduct the experiments, tuning, and sensitivity tests. Each experiment section in the following chapter highlight additions to this process necessary to perform the particular set of experiments.

Overarching Formulation Decision

A multi-option flexibility framework that pairs space logistics modeling via DES with decision rules and Monte Carlo simulations that sample multi-domain uncertainty will allow the user to draw conclusions on the relative performance of flexible option sets for LEO-based OOS




Figure 3.20: Use of the JAWA Framework




Figure 3.21: Experimentation Logic Flow Diagram




Figure 3.22: Experimental Method Diagram

CHAPTER 4

SCREENING FLEXIBLE OPTIONS & EXPERIMENTATION




Fig. 2—The “Bavaria”

Figure 4.1: The *Bavaria* of the Semmering Railway Contest [232]

”They built these tracks even before there was a train in existence that could make the trip. They built it because they knew some day, the train would come.”
- *Under the Tuscan Sun*

4.1 Experimentation Overview

The previous chapter developed the flexible option screening methodology by considering previous frameworks and determining which methodologies and tools best apply to the needs of OOS operations in LEO. The first hypothesis focuses on the collection-as-a-service concept without incremental deployment flexibility, determining which initial configuration performs best in terms of NOx emissions and cost, compared to the 0-warehouse baseline that relies exclusively on overpopulation for spare satellites and ADR vehicles launched from Earth to collect failed satellites. Figure 4.2 provides

a breakdown of all hypotheses and their experimental supporting conditions. If a configuration featuring CAAS outperforms the 0-warehouse baseline in the sustainability metrics for similar or better total cost, hypothesis 1 is substantiated and the resulting architecture forms the CAAS-baseline that is used for comparison in hypothesis 2. To determine which flexible options, or set of options, improves system performance in terms of cost and/or emissions, Experiment 2 compares flexible CAAS cases with both the inflexible CAAS-baseline and 0-warehouse baseline. To determine the sensitivity of the ranked configurations, the experiment includes a series of sensitivity studies on key uncertain variables and deterministic assumptions. Hypothesis 2 is substantiated if it identifies which set of flexible strategies and mechanisms bring the best economic and sustainability value. Lastly, Experiment 3 applies the effect of various policies to the costs and benefits and compares the results to those of the flexible, laissez-faire result from Experiment 2. The methodology incorporates government interventions through a parametric modeling approach, where regulatory mechanisms are represented as time-varying incentive and penalty structures. Eight distinct policy architectures are examined: fee-based systems with performance-linked reimbursements, conditional penalty frameworks, mandatory insurance requirements, and tax-based subsidy mechanisms for sustainable technology adoption. The evaluation framework employs weighted scoring methodologies to balance competing objectives across cost efficiency, environmental impact, and operational sustainability metrics. Performance indicators include total system costs, satellite collection and refurbishment rates, and atmospheric emission reductions. Experiment 3 is supported if it can identify configurations that improve sustainability metrics with manageable costs for both government and constellation operator over the 30-year scenario timeline. The results of experiments 1-3 are visualized and compared using Value at Risk/Value at Gain (VARG) plots for total cost, where value-at-risk represents the 95th percentile and value-at-gain represents the 5th percentile. To measure the statistical significance

of each configuration variable, this framework applies a polynomial regression-based variance decomposition approach with Analysis of Variance (ANOVA)-style F-tests to the VARG results. Configurations are compared by rank order across the various metrics. To determine the number of scenarios needed to determine if rank is stable, this thesis uses ranking convergence analysis that considers rolling statistics over sliding windows, monitoring how estimates evolve with additional scenarios. Coefficient of variation provides a normalized measure of rank variability while 95% confidence intervals provide insight into the typical range of a configuration's rank. These measurements help determine if rank is confidently stable for a given number of scenarios. Additionally, these experiments also directly compare the 5th percentiles, averages, and 95th percentiles of top-performing configurations to those of the baseline configurations. Using two-sample t-tests, the experiments determine if the p-value for the percent difference satisfies the standard $\alpha = 0.05$ threshold, if the 95% confidence interval includes zero, and whether the margin of error is smaller than the observed difference. In summary, rank order provides insight into which configurations perform best while VARG comparison informs the degree to which these configurations perform better.




Figure 4.2: Supporting Conditions for Hypotheses

4.1.1 Use Case

Borrowing the use case from Jakob et. al.'s multi-echelon sparing paper but accounting for recent updates, this thesis uses OneWeb as the example customer constellation. The use case resembles a portion of the constellation as it is today, with a constellation altitude of 1200 km, an inclination of 86.4 degrees, 18 planes with 36 satellites each, and an initial satellite failure rate of 1 satellite per year. CAAS infrastructures will start with at least 1 spare warehouse located at the optimal altitude, 796km, identified by Jakob et al. [14]. All configurations feature 2 in-plane spare satellites and 5 initial spare satellites per warehouse. If the servicing infrastructure grows to over 3 warehouses, each warehouse houses at least 2 spare satellites; if the servicing infrastructure grows to over 8 warehouses, this number is further reduced to 1 spare satellite. Note that the results for the following experiments are based on this use case and its assumptions. This use case provides a basis for demonstrating the proposed flexibility framework and its methodology, but the framework could be expanded in future work to consider other satellite constellations of varying parameters.

4.1.2 Testing for Statistical Significance

Individual Configuration Variables

The test for statistical significance methodology serves to identify which configuration variables have the most significant impact on key performance metrics in the simulation framework. This statistical approach helps the user understand the relative importance of different design decisions across complex scenarios involving multiple categorical configuration parameters like warehouse configurations, satellite configurations, and policy schemes.

The analysis begins by transforming categorical configuration variables into numerical form using one-hot encoding, which creates binary indicator variables for each

category level while avoiding multicollinearity issues. To capture complex relationships and interaction effects between different configuration choices, the methodology expands these encoded features using polynomial terms up to degree 3. This expansion includes individual configuration effects, pairwise interactions between configurations, and three-way interactions.

For each unique configuration combination, the system computes three key risk-adjusted performance metrics for total cost: average performance across simulation runs, Value at Risk (VaR) representing the 95th percentile of outcomes, and Value at Gain (VAG) representing the 5th percentile. These metrics provide a comprehensive view of both typical performance and extreme outcomes under uncertainty.

The analysis uses F-regression to assess how well each polynomial feature explains variance in the VARG metrics. The F-statistic measures the ratio of explained variance to unexplained variance for each feature, while p-values indicate statistical significance. The implementation handles various types of categorical variables, from binary choices like depot presence to multi-level variables like policy schemes. This process enables the user to prioritize configuration decisions by focusing on variables with the highest explanatory power, understand risk trade-offs by comparing variable importance across different metrics, identify synergistic relationships between configuration choices, and validate that simulation designs capture expected patterns of influence. The application across multiple performance metrics provides a comprehensive view of configuration sensitivity for evidence-based decision-making in complex space systems.

Configurations

The methodology addresses two different but complementary questions that arise when comparing multiple design configurations across stochastic scenarios. The first question asks how many uncertainty scenarios are needed before we can trust that

a configuration’s relative rank order is stable and the second question addresses the magnitude of performance differences. Stability in rank order ensures sufficient data collection before making comparisons, while significance testing provides the statistical rigor to justify specific quantitative claims about performance differences.

Ranking Convergence Analysis When running Monte Carlo simulations across multiple uncertainty scenarios, configuration rankings can vary considerably in early iterations as the simulation explores different regions of the uncertainty space. The ranking convergence analysis determines if there are sufficient scenarios to determine stable rank order. This analysis is essential because it balances computational expense with confidence in rank order.

The core methodology tracks each configuration’s rank across cumulative uncertainty scenarios and computes a rolling variance over a sliding window. For a sequence of ranks r_1, r_2, \dots, r_n observed across n cumulative scenarios, a sliding window of size w examines w consecutive rank values. The rolling variance at position i is computed as:

$$\sigma_{\text{rolling}}^2(i) = \frac{1}{w} \sum_{j=i}^{i+w-1} (r_j - \bar{r}_i)^2$$

where $\bar{r}_i = \frac{1}{w} \sum_{j=i}^{i+w-1} r_j$ is the mean rank within window i , and the window slides forward as each new scenario is added to the cumulative set. Unlike global variance, which computes a single statistic across all scenarios, rolling variance is a local, time-varying measure that tracks stability as scenarios accumulate. Each window provides a snapshot of rank variability over recent scenarios, enabling detection of convergence as the simulation progresses.

A configuration’s ranking is considered converged when all of three conditions are met simultaneously: the rolling variance must remain below 0.5 (rank units)² for at least three consecutive windows, the variance calculated over the most recent set of scenarios must also be below 0.5 (rank units)², and both conditions must be met at

the same time. The threshold of 0.5 is selected based on the discrete nature of ranks. Since ranks are dimensionless integers such as 1, 2, or 3, a variance of 0.5 (rank units)² indicates the rank is oscillating minimally, roughly staying within plus or minus one rank position. This level of variability represents acceptable stability for engineering decision-making while being stringent enough to avoid false convergence signals.

For each configuration, the analysis computes mean, median, mode, and standard deviation to characterize both central tendency and dispersion. The mean rank provides the average rank across all scenarios, offering insight into the configuration’s typical performance level. Bootstrap-derived 95% confidence intervals estimate the range within which the true underlying rank likely falls, accounting for sampling uncertainty.

Statistical Significance Testing Even when rankings have converged and relative orderings are stable, rigorous statistical methods are required to make specific quantitative claims about performance differences. While ranking convergence tells us that Configuration A consistently ranks higher than Configuration B, it does not quantify the difference between them. Statistical significance testing provides the quantitative rigor to support such claims with defensible confidence levels.

To compare two configurations, a two-sample independent t-test (Welch’s t-test) that does not assume equal variance is performed to determine whether the observed difference in means is statistically distinguishable from zero [233]. This calculation is done using `scipy stats.ttest_ind` [234]. The Welch t-test method calculates degrees of freedom using Equation 4.1. The null hypothesis posits that the true mean of the configuration equals the true mean of the baseline, implying no real difference exists between them. The alternative hypothesis states that the two means differ, indicating a genuine performance gap. The test statistic is calculated as the difference in sample means divided by the pooled standard error, which accounts for the variance and

sample size of both groups. Using a significance level of $\alpha = 0.05$, which represents the standard threshold, the condition is straightforward: if the computed p-value is less than 0.05, we reject the null hypothesis and conclude the difference is statistically significant; if the p-value is 0.05 or greater, we fail to reject the null hypothesis and conclude there is no statistically significant difference, meaning the configurations show comparable performance.

$$\nu = \frac{\left(\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}\right)^2}{\frac{\left(\frac{s_1^2}{n_1}\right)^2}{n_1-1} + \frac{\left(\frac{s_2^2}{n_2}\right)^2}{n_2-1}} \quad (4.1)$$

ν = degrees of freedom for Welch's t-test

s_1^2 = sample variance of group 1

s_2^2 = sample variance of group 2

n_1 = sample size of group 1

n_2 = sample size of group 2

Complementing hypothesis testing, 95% confidence intervals are computed for each configuration's performance metrics. For the mean performance, the confidence interval is calculated as the sample mean plus or minus the product of the critical t-distribution value and the standard error, where standard error equals the sample standard deviation divided by the square root of the sample size. For the 5th and 95th percentiles, confidence intervals are instead calculated using bootstrap resampling: 1,000 resampled datasets are generated by randomly sampling with replacement from the original data, the percentile is calculated for each resample, and the confidence interval bounds are determined as the 2.5th and 97.5th percentiles of these bootstrap estimates [235]. Bootstrap resampling is used for percentiles because, unlike means,

percentiles do not have simple mathematical formulas for their confidence intervals.

Regardless of the calculation method, confidence intervals provide three crucial pieces of information: the point estimate gives the best estimate of the true performance, the uncertainty bounds define the range within which the true value likely falls with 95% confidence, and the interval width serves as a precision indicator where narrow intervals indicate precise estimates based on low variability while wide intervals reveal high variability requiring caution in interpreting point estimates. The margin of error, defined as the half-width of the confidence interval expressed as a percentage of the point estimate, provides an intuitive measure of estimate precision across all statistics.

Integration of the Two Methodologies The typical workflow proceeds in stages: first, ranking convergence analysis is performed after initial scenario batches to assess whether configurations have achieved stable relative orderings. Configurations with converged rankings proceed to the significance testing phase, while those exhibiting low confidence or lack of convergence indicate that they require additional scenario collection. Once convergence is confirmed, significance testing evaluates the magnitude of observed differences, distinguishing between genuine performance gaps and statistical noise. The combination enables confident statements of the form: "Configuration A consistently ranks first across scenarios (ranking convergence), and achieves a statistically significant 5% cost reduction compared to baseline (significance testing)."

Ranking convergence analysis serves as a prerequisite quality check, ensuring that sufficient scenarios have been executed before significance testing is conducted. The ranking convergence analysis ensures computational efficiency by identifying when additional scenarios provide negligible new information, while the significance testing provides the evidential foundation needed to support specific performance claims.

4.1.3 Sensitivity Testing

The results of this framework provide preliminary, rank-ordered solutions to provide guidance on which investments could encourage OOS in LEO. Global and local sensitivity studies can verify assumptions and distributions in the framework. Lin demonstrates this practice in their thesis, subjecting their simulation to a series of tests to check their assumptions and highlight the best candidates for future work [208]. Lin toggles on reservoir, facility, and market uncertainty to conduct their sensitivity study. By comparing the impact of different uncertainties, they determine which ones influence the shape of the VARG curves. Highly influential variables are good candidates for future refinement. Variables with low influence, meanwhile, are sufficiently represented with the provided assumptions and distributions. Lin also conducts sensitivity studies on assumed, deterministic values such as the cost of options, benefit of options, and implementation time, first parameterizing nominal values over a wide range and observing the change in rank order. Subsequently, Lin makes small perturbations for their local sensitivity study, observing the change of the economics statistics. Checking sensitivity of distributions and assumptions provides a form of validation for the methodology choices and highlights opportunities for future work.

To measure the effect of key sources of uncertainty on the results, this thesis evaluates the impact of various aspects and sources of uncertainty, such as:

- **Static launch and satellite manufacturing cost:** rather than varying these parameters over time, this test holds launch and satellite manufacturing cost constant over the 30-year duration
- **Removed Effect of Technology Obsolescence:** The simulated time to obsolescence in this sensitivity test is set to be greater than the simulated duration
- **Varying Initial Failure Rate:** The initial failure rate (1 failure per year)

is set to 0.75 and 0.5 failures per year to understand the relationship between satellite reliability and relative CAAS performance

- **ADR Vehicle Deployment:** This sensitivity test only deploys ADR vehicles by launching them from Earth; the CAAS warehouses do not deploy ADR vehicles
- **Reusable Second Stage Vehicles:** This sensitivity test does not incorporate reusable second stage vehicles

Perturbation Analysis for Deterministic Assumptions

Perturbation analysis provides a method for evaluating the robustness of system configuration rankings under parametric uncertainty. This approach addresses the critical challenge of determining whether apparently superior configurations maintain their performance advantages when estimated input parameters deviate from their initially assumed values. By varying uncertain cost parameters and comparing the resulting configuration rank order, this approach allows the user to determine which configuration settings are robust. It also highlights which assumed values would require greater attention and a higher degree of confidence in future work.

Rather than perturb every assumed/estimated deterministic input, the perturbation analysis focuses on a subset of inputs based on cost frequency and magnitude. This selection approach ensures that perturbation analysis captures economically consequential uncertainties that could significantly affect long-term system performance while establishing a manageable and computationally efficient process.

Key parameters selected for perturbation analysis include:

- **ADR vehicle launch and deployment cost:** major operational expense affecting service delivery economics

- **Satellite refurbishment cost:** high-frequency operational expense directly impacting service value proposition
- **Satellite repair cost:** frequent operational cost affecting maintenance strategy decisions
- **ADR operation cost:** recurring service cost influencing collection economics
- **Satellite refueling cost:** operational expense affecting satellite lifetime extension strategies
- **Downmass return cost:** expense of returning satellites back to Earth on a reusable second stage vehicle
- **ADR cost min fraction:** The minimum fraction of present-day ADR vehicle costs that is achievable within the learning-curve ADR cost model

The variables listed here are perturbed by X1.25 and X1.5 or X0.5 and X0.25. These perturbation magnitudes reflect a compromise between capturing meaningful uncertainty ranges while maintaining computational tractability. This level represents a realistic assessment of parameter uncertainty in many engineering applications, encompassing forecasting errors, technological variations, and market volatility.

The analysis evaluates how configuration rankings change under perturbations, providing insight into the reliability of configuration selection decisions. Configurations that maintain their relative performance advantages across perturbation scenarios demonstrate robustness, while those exhibiting high sensitivity to parameter changes may represent risk if the assumed input lacks a high degree of confidence.

4.2 H1 Experimentation: Collection as a Service

The first set of experiments focuses on the economic and environmental performance of different options in a laissez-faire environment, without flexible deployment deci-

sions and without the simulated effect of market-based policy. To review, the research questions and related hypothesis are as follows:

Research Question 1: Economic Value

Which collection hub infrastructure configuration provides the greatest economic value for OOS providers compared to traditional practices?

Research Question 2: Emission Reduction

Which collection hub infrastructure configuration provides the greatest reduction in atmospheric emissions compared to traditional overpopulation sparing strategies?

Hypothesis 1

Multi-echelon sparing with collection warehouses that strategically retain salvageable assets for future reuse, refueling, or repair, while facilitating ADR missions and the controlled deorbiting or Earth-return of old satellites via reusable second stages provide better or comparable economic feasibility and sustainability metrics compared to infrastructures that rely on the overpopulation sparing strategy and ADR missions launched directly from Earth.




Figure 4.3: Experiment 1 Methodology to Support Hypothesis 1

Experiment 1 provides equivalence to the real-world problem by considering a reality where satellite operators are obligated to clean up failed satellites in their constellations. The experiment reflects realistic factors including uncertain technology

costs (launch and manufacturing), stochastic satellite failures, and the economic pressures that currently favor the make-use-dispose approach to constellation operations. This experiment compares three strategies, starting with the 0-warehouse baseline, where the operator uses overpopulation to replenish constellations after failures and launches ADR vehicles from Earth every time there is a satellite failure. The second strategy is the multi-echelon sparing strategy that uses parking spares and in-plane spares to reduce costs associated with replenishing the satellite constellation. This strategy also deploys ADR vehicles exclusively from Earth. Lastly, the CAAS strategy builds upon the multi-echelon parking strategy to provide collection services as well as ADR docking and refueling. The experiment considers various versions of the future with uncertain variables sampled from distributions or a DOE. The anticipated decrease in launch cost and spacecraft manufacturing costs are critical sources of uncertainty, since they make it easier to continue using the make-use-dispose approach to satellite constellation operations. These scenarios do not predict the future so much as model the effect of various futures to highlight which aspects of the CAAS architecture, if any, provide consistent and measurable benefits compared to the other baseline strategies. In doing so, this experiment addresses gaps from literature, namely, which aspects of CAAS provide improvements in cost and emissions and which aspects are worthy of continued study. This experiment highlights opportunities to improve circularity in LEO-based space operations and determines if, to the extent that the CAAS infrastructure is able to service satellites, it provides measurable benefits, even if servicing is not extended for the entire constellation. This experiment seeks to understand which technology investments, such as RPO-capable satellites, servicable satellites, or collection warehouses capable of refueling/repairing satellites in space, provide the most compelling case for the overall CAAS concept.

Table 4.1: Space Infrastructure Configuration Parameters and Dependencies

Parameter	Values	Dependency	Description
Number of warehouses	$\{0,1,2,3,4\}$	None	Depot infrastructure configuration: 0 = No warehouse (baseline); 1–3 = Number of warehouses.
CAAS Mode	$\{0,1\}$	$\# \text{ warehouses} \neq 0$	Collection-as-a-Service availability: 0 = Disabled; 1 = Enabled.
Refuelable/Repairable Upgrades	$\{0,2\}$	$\# \text{ warehouses} \neq 0$ and CAAS = 1	Refuelable satellites: 0 = No refuelable capability; 2 = Refuelable capability.
Warehouse Upgrades	$\{0,2\}$	Upgraded satellites = 2	Warehouse upgrade strategy: 0 = No upgrades; 2 = Immediate upgrades.
RPO-Capable Satellites	$\{0,2\}$	$\# \text{ warehouses} \neq 0$ and CAAS = 1	Rendezvous and Proximity Operations capability: 0 = Disabled; 2 = Available.

Table 4.2: Configuration Combination Summary

Configuration Type	Count	Key Characteristics
Baseline (No Depot)	1	All parameters = 0
Depot without CAAS	3	$\text{depot_config} \in \{1,2,3,4\}$, $\text{CAAS_config} = 0$
CAAS without Refuelable	8	$\text{CAAS_config} = 1$, $\text{refuelable_config} = 0$, $\text{rpo_config} \in \{0,2\}$
Full CAAS System	16	$\text{CAAS_config} = 1$, $\text{refuelable_config} = 2$, $\text{upgraded_config} \in \{0,2\}$, $\text{rpo_config} \in \{0,2\}$
Total Combinations	29	Per scenario iteration

4.2.1 Configuration Set

These configuration settings and dependencies produce 29 unique configurations, described in Table 4.8, where `CAAS_config` is CAAS mode, `depot_config` is the number of warehouses, `refuelable_config` (Rf) is whether satellites are upgraded to be refuelable and repairable, `rpo_config` (RPO) is whether satellites are RPO-capable, and `upgraded_config` (Upg) is whether warehouses are upgraded to service satellites.

4.2.2 VARG Results for Experiment 1

This section presents the results of the Value at Risk and Gain (VARG) analysis combined with ranking convergence assessment and statistical significance testing for multiple performance metrics across different space operations configurations. The analysis evaluated 29 unique configurations through 80 uncertainty scenarios across seven key performance metrics to identify superior configurations, assess ranking stability, and determine the statistical significance of observed performance differences.

The output metrics are:

- Customer costs: the total cost that the satellite constellation operator accrues over the 30 year timeline, including costs related to operating the constellation as well as the CAAS infrastructure
- Total Operational costs: the total cost that the satellite constellation operator accrues over the 30 year timeline, excluding the initial infrastructure cost
- Total Emissions: the total NOx emissions, in kg, from all spacecraft reentering the atmosphere, accounting for both intact reentry and demise
- Total refurbished satellites: the total number of satellites that have been refurbished, either on Earth or in space

- Satellite refurbished in space: the total number of satellites that have been refurbished in space
- Total Cost over Total Collections: the total number of collected satellites compared to the total cost of the entire system
- Total Cost over Total Refurbishments: the total number of refurbished satellites compared to the total cost of the entire system

The VARG analysis computes average performance, Value at Risk and Value at Gain for each metric across all tested configurations. The ranking convergence analysis confirmed that 80 uncertainty scenarios provide stable rank orderings for key performance metrics, with the top-performing configurations achieving convergence confidence levels. This ranking stability in top performing configurations across 80 scenarios provides confidence that the relative performance orderings identified represent robust patterns rather than transient fluctuations in early scenario iterations.

Experiment 1 Results

The analysis of configurations spanning zero to four warehouses reveals distinct patterns in both relative ranking and absolute performance differences. The top ranking configuration for each output is provided in Table 4.3. The ranking convergence analysis confirms that the optimal CAAS configuration for total cost (D1, C1, Rf0, Up0, RP2) maintains consistent second-place ranking across uncertainty scenarios, demonstrating that a single-depot CAAS configuration with RPO-capable satellites and without satellite or warehouse upgrades reliably provides the most cost-effective CAAS solution. While the 0-warehouse baseline configuration (D0, C0, Rf0, Up0, RP0) has a stable 1st place ranking, when examining the magnitude of this advantage through statistical significance testing, the results reveal a more nuanced picture of cost comparability.

Table 4.3: Parameter Settings for Top Performing Configurations: Experiment 1

Configuration	Depot	Upgraded	Refuelable	RPO	CAAS
0-Warehouse Baseline	0	0	0	0	0
Best Cost No-CAAS	1	0	0	0	0
Best for Total Cost & Weighted Multi-Objective	1	0	0	2	1
Best for Total Emissions & Total Num Sat Refurb. & Cost/-Collection	4	0	2	2	1
Best for Cost/Refurbish & Sats Refurb. In Space	4	2	2	2	1
Best for Op Costs	4	0	0	2	1

The cost-superior CAAS configuration, shown in Table 4.4, achieves an average total cost of \$4.49B, representing a 0.77% increase compared to the 0-warehouse baseline of \$4.45B. Statistical significance testing indicates this difference is not statistically distinguishable [$p=0.59$, 95% CI: [4.40e9, 4.57e9], MoE: 1.90%], demonstrating comparable cost performance despite the stable superior ranking. This finding illustrates an important distinction between ranking stability and statistical significance: the single-depot CAAS configuration consistently ranks second across scenarios, yet the magnitude of its cost disadvantage is modest enough that it falls within the uncertainty bounds of the baseline performance. Decision-makers can therefore conclude that CAAS reliably offers cost-comparable performance to the 0-warehouse baseline. This similarity in total cost is illustrated in the total cost VARG plot, provided in Figure 4.4. Table 4.4 includes the top-ranked configuration for each output metric and compares their total cost across percentiles. This table also includes their 95% confidence intervals, margin of error, percent comparison to the 0-warehouse baseline and multi-echelon baseline and the respective p-values.

The emissions performance, illustrated in Figure 4.5 and tabulated in Table 4.5, presents a contrasting pattern where both ranking superiority and statistical signifi-





Figure 4.4: Total Cost for Experiment 1: VARG Plot

cance align to demonstrate clear CAAS advantages. The emissions-superior configuration (D4, C1, Rf2, Up0, RP2) with four depots and serviceable satellites achieves -10.58% lower average emissions compared to the baseline ($p < 0.001$, 95% CI: [3.04M kg, 3.20M kg], MoE = 2.41%), representing a statistically significant and substantial reduction. Even the cost-optimal CAAS configuration, which prioritizes economic efficiency over environmental performance, achieves -6.84% lower emissions on average ($p < 0.001$, 95% CI: [3.17M kg, 3.33M kg], MoE = 2.45%), a reduction that is both statistically significant and operationally meaningful. These emissions advantages persist across all performance percentiles: at the 5th percentile, CAAS configurations achieve -9.75% (top cost configuration) and -9.69% (top emissions configuration) reductions, while even at the 95th percentile, emissions remain -5.44% and -8.67% lower than baseline for the top cost and top emissions configurations, respectively. The consis-

tency of emissions benefits across the entire performance distribution, combined with strong statistical significance, demonstrates that environmental advantages represent fundamental architectural benefits of the warehouse approach rather than outcomes contingent on favorable conditions.



The comparison with the multi-echelon sparing baseline without CAAS provides additional context for evaluating the CAAS value proposition. The best no-CAAS multi-echelon configuration (D1, C0, Rf0, Up0, RP0, P0) has 3.29% higher costs than the 0-warehouse baseline, a difference that is statistically significant ($p = 0.025$, 95% CI: [\$4.51B, \$4.69B], MoE = 1.97%). Additionally, this multi-echelon configuration shows no statistically significant emissions improvement over the baseline (0.22% difference, $p = 0.89$, 95% CI: [3.42M kg, 3.58M kg], MoE = 2.24%), indicating that passive depot infrastructure without active collection capability or ADR deployment

provides no environmental benefits. This comparison validates that the CAAS operational model specifically, not merely the presence of depot infrastructure for sparing purposes, drives both the cost feasibility and environmental advantages observed in the analysis.

The refurbishment throughput metric, shown in Table 4.6, reveal the fundamental architectural difference between CAAS and traditional sparing strategies. The cost-optimal CAAS configuration enables an average of 65 satellite refurbishments over the 30-year simulation period with 95% CI [54,75]. The emissions-optimal configuration with greater infrastructure investment achieves 191 average refurbishments with 95% CI [169, 213]. In contrast, both the 0-warehouse baseline and multi-echelon configurations without CAAS achieve zero refurbishments by definition. While the absolute refurbishment numbers remain modest relative to the scale of mega-constellation operations, they represent meaningful progress toward circular space systems and demonstrate proof-of-concept for satellite lifetime extension through collection and servicing infrastructure. They highlight that an OOS system does not need to service an entire constellation to provide meaningful benefits.

Refurbishment rates carry multiple sustainability implications: fewer satellites burning up in the atmosphere translates to decreased manufacturing demand for replacement satellites, and lower launch requirements with their associated environmental footprints. Each satellite refurbished rather than replaced represents avoided lifecycle environmental impacts from material extraction, component manufacturing, assembly operations, and launch services. While this framework focuses primarily on atmospheric NOx emissions, refurbishment rates indicate broader circular economy benefits that extend beyond the explicitly modeled environmental metrics.

Table 4.4: Experiment 1: Total Cost

Top Performing Metric	Configuration	Percentile	Total Cost Value (\$)	Total Cost CI Lower	Total Cost CI Upper	Total Cost MoE%	Total Cost % vs Baseline	Total Cost p-value (Baseline)	Total Cost Sig? (Baseline)	Total Cost % vs No-CAAS	Total Cost p-value (No-CAAS)	Total Cost Sig? (No-CAAS)
0-Warehouse Baseline	d:0, up:0, rf:0, RPO:0, CAAS:0	avg	4.45E+09	4.36E+09	4.54E+09	2.085	N/A	N/A	N/A	-3.188	0.026	Yes
		5th	3.95E+09	3.89E+09	4.04E+09	1.906	N/A	N/A	N/A	-3.981	0.026	Yes
		95th	5.25E+09	5.01E+09	5.69E+09	6.443	N/A	N/A	N/A	-3.146	0.026	Yes
Best Cost No-CAAS	d:1, up:0, rf:0, RPO:0, CAAS:0	avg	4.60E+09	4.51E+09	4.69E+09	1.972	3.293	0.026	Yes	N/A	N/A	N/A
		5th	4.12E+09	4.10E+09	4.21E+09	1.321	4.146	0.026	Yes	N/A	N/A	N/A
		95th	5.43E+09	5.09E+09	5.82E+09	6.728	3.249	0.026	Yes	N/A	N/A	N/A
Best Total Cost & Weighted Multi-Objective	d:1, up:0, rf:0, RPO:2, CAAS:1	avg	4.49E+09	4.40E+09	4.57E+09	1.896	0.766	0.590	No	-2.446	0.074	No
		5th	4.06E+09	4.03E+09	4.10E+09	0.841	2.702	0.590	No	-1.387	0.074	No
		95th	5.18E+09	4.98E+09	5.56E+09	5.582	-1.474	0.590	No	-4.574	0.074	No
Best emissions, Total Num Sat Refurb, cost/collect	d:4, up:0, rf:2, RPO:2, CAAS:1	avg	4.86E+09	4.78E+09	4.94E+09	1.685	9.195	0.000	Yes	5.713	0.000	Yes
		5th	4.44E+09	4.39E+09	4.49E+09	1.122	12.194	0.000	Yes	7.728	0.000	Yes
		95th	5.54E+09	5.27E+09	5.81E+09	4.846	5.525	0.000	Yes	2.205	0.000	Yes
Best for Cost/Refurbish & Sats Refurb. In Space	d:4, up:2, rf:2, RPO:2, CAAS:1	avg	5.01E+09	4.92E+09	5.09E+09	1.673	12.507	0.000	Yes	8.920	0.000	Yes
		5th	4.56E+09	4.51E+09	4.62E+09	1.206	15.350	0.000	Yes	10.758	0.000	Yes
		95th	5.73E+09	5.50E+09	5.99E+09	4.330	9.141	0.000	Yes	5.707	0.000	Yes
Best for Operational Costs	d:4, up:0, rf:0, RPO:2, CAAS:1	avg	4.84E+09	4.76E+09	4.92E+09	1.674	8.681	0.000	Yes	5.216	0.000	Yes
		5th	4.42E+09	4.36E+09	4.45E+09	1.018	11.801	0.000	Yes	7.350	0.000	Yes
		95th	5.47E+09	5.27E+09	5.77E+09	4.521	4.063	0.000	Yes	0.789	0.000	Yes

Table 4.5: Experiment 1: Total NOx Emissions

Top Performing Metric	Configuration	Percentile	Total Emissions Value(kg)	Total Emissions CI Lower	Total Emissions CI Upper	Total Emissions MoE%	Total Emissions % vs Baseline	Total Emissions p-value (Baseline)	Total Emissions Sig? (Baseline)	Total Emissions % vs No-CAAS	Total Emissions p-value (No-CAAS)	Total Emissions Sig? (No-CAAS)
0-Warehouse Baseline	d:0, up:0, rf:0, RPO:0, CAAS:0	avg	3.49E+06	3.41E+06	3.56E+06	2.165	N/A	N/A	N/A	-0.223	0.887	No
		5th	2.88E+06	2.74E+06	3.02E+06	4.786	N/A	N/A	N/A	4.035	0.887	No
		95th	3.95E+06	3.88E+06	4.12E+06	3.047	N/A	N/A	N/A	-0.778	0.887	No
Best Cost No-CAAS	d:1, up:0, rf:0, RPO:0, CAAS:0	avg	3.50E+06	3.42E+06	3.58E+06	2.240	0.223	0.887	No	N/A	N/A	N/A
		5th	2.77E+06	2.72E+06	3.04E+06	5.897	-3.879	0.887	No	N/A	N/A	N/A
		95th	3.98E+06	3.89E+06	4.07E+06	2.179	0.784	0.887	No	N/A	N/A	N/A
Best Total Cost & Weighted Multi-Objective	d:1, up:0, rf:0, RPO:2, CAAS:1	avg	3.25E+06	3.17E+06	3.33E+06	2.455	-6.844	0.000	Yes	-7.051	0.000	Yes
		5th	2.60E+06	2.53E+06	2.79E+06	4.926	-9.747	0.000	Yes	-6.105	0.000	Yes
		95th	3.73E+06	3.67E+06	3.84E+06	2.800	-5.440	0.000	Yes	-6.176	0.000	Yes
Best emissions, Total Num Sat Refurb, cost/collect	d:4, up:0, rf:2, RPO:2, CAAS:1	avg	3.12E+06	3.04E+06	3.20E+06	2.416	-10.582	0.000	Yes	-10.782	0.000	Yes
		5th	2.60E+06	2.39E+06	2.64E+06	4.834	-9.689	0.000	Yes	-6.044	0.000	Yes
		95th	3.60E+06	3.49E+06	3.76E+06	3.770	-8.673	0.000	Yes	-9.383	0.000	Yes
Best for Cost/Refurbish & Sats Refurb. In Space	d:4, up:2, rf:2, RPO:2, CAAS:1	avg	3.15E+06	3.07E+06	3.23E+06	2.480	-9.805	0.000	Yes	-10.006	0.000	Yes
		5th	2.54E+06	2.43E+06	2.65E+06	4.349	-11.678	0.000	Yes	-8.114	0.000	Yes
		95th	3.65E+06	3.56E+06	3.75E+06	2.649	-7.605	0.000	Yes	-8.323	0.000	Yes
Best for Operational Costs	d:4, up:0, rf:0, RPO:2, CAAS:1	avg	3.21E+06	3.14E+06	3.28E+06	2.201	-8.058	0.000	Yes	-8.262	0.000	Yes
		5th	2.67E+06	2.51E+06	2.81E+06	5.569	-7.144	0.000	Yes	-3.397	0.000	Yes
		95th	3.66E+06	3.56E+06	3.80E+06	3.263	-7.165	0.000	Yes	-7.887	0.000	Yes

Table 4.6: Experiment 1: Number of Refurbishments

Top Performing Metric	Configuration	Percentile	Num Sat Refurb Value	Num Sat Refurb CI Lower	Num Sat Refurb CI Upper	Num Sat Refurb MoE%
0-Warehouse Baseline	d:0, up:0, rf:0, RPO:0, CAAS:0	avg	0.00	0.00	0.00	N/A
		5th	0.00	0.00	0.00	N/A
		95th	0.00	0.00	0.00	N/A
Best Cost No-CAAS	d:1, up:0, rf:0, RPO:0, CAAS:0	avg	0.00	0.00	0.00	N/A
		5th	0.00	0.00	0.00	N/A
		95th	0.00	0.00	0.00	N/A
Best Total Cost & Weighted Multi-Objective	d:1, up:0, rf:0, RPO:2, CAAS:1	avg	64.91	54.11	75.71	16.636
		5th	0.00	0.00	0.00	-
		95th	143.20	121.00	184.00	21.997
Best emissions, Total Num Sat Refurb, cost/collect	d:4, up:0, rf:2, RPO:2, CAAS:1	avg	191.40	169.72	213.08	11.328
		5th	27.55	0.00	59.90	108.711
		95th	319.25	304.20	345.00	6.390
Best for Cost/Refurbish & Sats Refurb. In Space	d:4, up:2, rf:2, RPO:2, CAAS:1	avg	184.85	171.71	197.99	7.108
		5th	100.40	83.00	114.75	15.812
		95th	282.30	260.00	310.00	8.856
Best for Operational Costs	d:4, up:0, rf:0, RPO:2, CAAS:1	avg	157.98	140.32	175.63	11.176
		5th	3.80	0.00	49.30	648.684
		95th	260.75	243.45	287.00	8.351

Operational cost analysis provides insight into the economic mechanisms underlying CAAS performance. Excluding initial capital expenditure and focusing solely on ongoing operational expenses, the cost-optimal CAAS configuration achieves - 2.81% lower operational costs compared to the 0-warehouse baseline ($p = 0.078$, 95% CI: [\$3.22B, \$3.39B], MoE = 2.57%). While this operational cost advantage shows marginal significance ($p = 0.078$, slightly above the $\alpha = 0.05$ threshold), the trend suggests genuine operational efficiencies from asset reuse and lifetime extension. The near-significance combined with the narrow confidence interval indicates that with modest additional scenarios, this operational advantage would likely achieve full statistical confirmation. More importantly, the operational cost reduction quantifies the ongoing economic benefits that offset capital investment, identifying the necessary budget range to make CAAS costs equal to the 0-warehouse baseline.

Experimental Observation 1.1

CAAS demonstrates statistically significant emissions reductions and refurbishment rates compared to the 0-warehouse baseline, while achieving cost-comparable performance. Multi-echelon sparing without CAAS shows neither cost advantages nor emissions benefits, validating that active collection capability specifically drives the observed sustainability improvements and cost feasibility.

Constraining the design space to a maximum of 1 depot, the analysis focuses comparison between the most economically viable CAAS architectures and the 0-warehouse baseline. In this case, the emissions-superior single-depot configuration is the same as the cost-superior configuration (D1, C1, Rf0, Up0, RP2, P0). Predictably, the refurbishment-superior single-depot configuration features both upgraded satellites and warehouses (D1, C1, Rf2, Up2, RP2, P0), but its improvement in refurbishments over the top cost/emission configuration is not statistically significant.

The emissions reduction mechanisms identified through the analysis reveal important architectural insights about CAAS sustainability benefits. Upgraded warehouses (enabling in-space satellite servicing) do not demonstrate statistical significance for emissions reduction, as configurations with and without this capability show comparable emissions performance when other parameters are held constant. Instead, the primary emissions drivers emerge from the number of warehouses, Earth-return dual-mission operations, and reusable second-stage adoption. When launch vehicles conduct combined missions, such as resupplying orbital warehouses or replenishing constellation satellites before collecting the old satellites for Earth-return, the consolidated operations reduce overall launch requirements and associated emissions. Similarly, the transition toward reusable second stages that return to Earth rather than burning up in the atmosphere provides substantial emissions benefits as the launch vehicle fleet modernizes over the simulation period. The ability to prevent satellites from burning up through controlled deorbiting contributes to emissions reduction, but the more substantial benefit comes from reducing launch cadence requirements and launch vehicle demise through asset reuse.

The refuelable satellite parameter demonstrates complex interactions with emissions outcomes. While refuelable satellites enable in-space servicing that extends satellite lifetimes, the servicing operations themselves require frequent propellant delivery launches that can increase emissions if not carefully managed. Additionally, refuelable and modular satellites provide benefits for Earth-based refurbishment by streamlining the process, which is represented by an uncertain factor that reduces Earth-based refurbishment costs. This highlights a critical consideration for sustainable space infrastructure design: well-intentioned servicing capabilities must be evaluated holistically to ensure they do not inadvertently exacerbate environmental impacts through increased launch activity. The current framework's focus on NOx emissions from atmospheric reentry represents a conservative environmental account-

ing that omits numerous lifecycle impacts from satellite manufacturing, launch vehicle production, and ground operations. A comprehensive lifecycle analysis incorporating multiple emission types and manufacturing impacts would provide a more complete environmental characterization, though the current NOx-focused analysis combined with the refurbishment metric offers reasonable proxies for broader sustainability performance.

4.2.3 Experimental Support for Hypothesis 1

Hypothesis 1 posits that multi-echelon sparing with collection warehouses that strategically retain salvageable assets for future reuse, refueling, or repair, while facilitating ADR missions and the controlled deorbiting or Earth-return of old satellites via reusable second stages provide better or comparable economic feasibility and improved sustainability metrics compared to infrastructures that rely on the overpopulation sparing strategy and ADR missions launched directly from Earth. The experimental evidence from both ranking convergence analysis and statistical significance testing provides comprehensive support for these claims, though with important nuances in the interpretation of "better or comparable" economic performance.

Economic Feasibility: Comparable Performance with Consistent Advantages

The economic feasibility claim receives strong but nuanced support from the dual statistical methodology. Ranking convergence analysis confirms that the cost-optimal CAAS configuration (single depot with RPO-capable satellites) maintains stable second-place ranking across the 80 uncertainty scenarios with the 0-warehouse baseline taking first place, demonstrating consistent relative superiority in cost performance relative to other CAAS configurations. Statistical significance testing reveals that this ranking advantage reflects modest absolute differences rather than dramatic cost improvements compared to the 0-warehouse baseline and other 1-depot CAAS

configurations, establishing cost parity rather than superiority for average performance outcomes.

The cost-superior CAAS configuration achieves an average total cost of \$4.49B, representing a 0.77% increase compared to the 0-warehouse baseline of \$4.45B. Statistical significance testing indicates this difference is not statistically distinguishable [$p=0.59$, 95% CI: [4.40e9, 4.57e9], MoE: 1.90%], demonstrating comparable cost performance despite the stable superior ranking. This combination of stable ranking with statistically insignificant absolute differences represents a valuable finding rather than a limitation. It demonstrates that CAAS reliably achieves the "comparable economic feasibility" condition specified in the hypothesis. The hypothesis explicitly includes "comparable" alongside "better" economic performance, recognizing that cost parity combined with sustainability advantages represents a viable value proposition, and the experimental evidence confirms this expectation.

The comparison with multi-echelon sparing without CAAS capability provides critical context for evaluating the warehouse value proposition. The best no-CAAS depot configuration incurs 3.29% higher costs than baseline ($p = 0.025$, statistically significant), demonstrating that passive depot infrastructure without active collection capability and ADR deployment provides no economic benefit. This establishes that the CAAS operational model specifically enables the cost competitiveness observed in the analysis.

Operational cost analysis provides insight into the economic mechanisms underlying CAAS performance. Excluding initial capital expenditure and focusing solely on ongoing operational expenses, the cost-optimal CAAS configuration achieves -2.81% lower operational costs compared to baseline ($p = 0.078$, 95% CI: [\$3.22B, \$3.39B], MoE = 2.57%). While this operational cost advantage shows marginal significance ($p = 0.078$, just above the $\alpha = 0.05$ threshold), the trend suggests genuine operational efficiencies from the CAAS concept. More importantly, the operational cost reduction

quantifies the ongoing economic benefits that offset capital investment, identifying the necessary budget range to make CAAS costs equal to the 0-warehouse baseline.

Sustainability Performance: Notable Improvements

The sustainability claims receive unambiguous and comprehensive empirical support from both ranking and significance analyses. The emissions metric demonstrates both stable superior ranking and statistically significant absolute advantages across all CAAS configurations and performance percentiles. The cost-optimal CAAS configuration achieves -6.84% lower average emissions compared to baseline ($p < 0.001$, 95% CI: [3.17M kg, 3.33M kg], MoE = 2.45%), a reduction that is both statistically significant and operationally meaningful in absolute terms. The emissions-optimal configuration with greater infrastructure investment delivers even stronger performance at -10.58% ($p < 0.001$, 95% CI: [3.04M kg, 3.20M kg], MoE = 2.41%), though at higher (and statistically significant) cost. These emissions advantages persist across percentiles, indicating fundamental architectural benefits rather than scenario-dependent outcomes.

The comparison with multi-echelon sparing without CAAS proves particularly revealing for understanding sustainability mechanisms. The best no-CAAS depot configuration shows no statistically significant emissions improvement over the baseline, confirming that passive depot infrastructure alone provides negligible environmental benefits. Combined with the significant cost penalty for no-CAAS configurations, this establishes that active collection and ADR-deployment capability specifically drives both the cost competitiveness and environmental advantages observed for CAAS. The hypothesis's architectural prescription of collection warehouses with strategic asset retention and ADR deployment are essential for achieving the claimed benefits.

Mechanism Validation and Architectural Insights

The RPO capability (satellite rendezvous and proximity operations) proves to be consistently beneficial across multiple metrics, enabling satellites to maneuver to collection locations and facilitating efficient servicing operations. This capability maintains stable ranking for both cost and emissions outcomes, validating the hypothesis' emphasis on collaborative maneuvering as a core architectural element. Meanwhile, the warehouse upgrade capability (enabling in-space satellite servicing) does not demonstrate statistical significance for emissions reduction when considered in isolation, contrary to initial expectations. Configurations with and without this capability show comparable emissions when other parameters are held constant, indicating that Earth-return refurbishment may offer better environmental and economic trade-offs than in-space servicing under current technology and cost assumptions. This trend is evident in Figure 4.6, which shows total NOx emissions compared to the year that Earth-return capability matures. When the capability becomes available earlier in the timeline, there is more opportunity to reduce emissions, as evidenced by the wide spread on the left hand side of the plot. As it takes longer for Earth-return to become available, there is a higher average total emissions.

The refuelable satellite parameter exhibits complex interactions requiring careful interpretation. While refuelable satellites enable in-space servicing that extends satellite lifetimes, the servicing operations themselves require frequent propellant delivery launches that can increase emissions if not carefully managed. Serviceable satellites appear in the top emissions and top refurbishments configurations (along with 4 depots), but when the design space is limited to 1 depot, the top emissions configuration (D1, C1, Rf0, Up0, RP2) does not feature serviceable satellites. This suggests serviceable satellites make a larger impact on emissions when coupled with larger CAAS infrastructures. The configuration (D4, C1, Rf2, Up0, RP2) achieves -10.58% emissions compared to the 0-warehouse baseline while the same configuration




Figure 4.6: NOx Emissions vs. Year of Earth-Return Maturity

without serviceable satellites has -8.06% emissions. Critically, this provides a larger difference than the comparison between (D1, C1, Rf2, Up0, RP2) and (D1, C1, Rf0, Up0, RP2). This difference is notable but not measurable with statistical significance for the given number of scenarios.

Limitations and Scope Considerations

Several limitations constrain the interpretation of results and suggest areas for future investigation. The economic model assumptions regarding CAAS capital costs significantly influence cost parity conclusions. If infrastructure costs decrease through technological learning or if satellite replacement costs increase due to capability enhancements, the economic case could shift from parity toward clear superiority. The analysis may not fully capture the economic value of extended satellite operational

life, particularly for high-value payloads where service continuity could provide benefits through enhanced mission flexibility.

The environmental modeling remains limited to NOx emissions from atmospheric reentry, omitting numerous lifecycle impacts from satellite manufacturing, launch vehicle production, material extraction, and ground operations. A comprehensive lifecycle analysis incorporating multiple emission types, manufacturing impacts, and end-of-life disposal options would provide more complete environmental characterization. The current framework also does not fully model the technical risks and operational complexities of rendezvous, docking, refueling, and component replacement operations that could affect real-world cost and reliability, potentially underestimating implementation challenges.

Market structure assumptions based on single-operator models may not capture network effects, risk pooling, or economies of scale that could emerge in a multi-operator CAAS marketplace serving diverse constellation operators. If multiple constellations share warehouse infrastructure and services, fixed costs could be amortized across larger customer bases, potentially transforming the economic case from marginal cost parity to clear cost advantages. The 30-year simulation timeframe may not fully capture long-term trends in launch cost reduction, satellite capability enhancement, or regulatory evolution that could influence CAAS value propositions over extended periods.

4.2.4 Hypothesis 1 Substantiation

The experimental evidence validates the central claims in Hypothesis 1 through the combined ranking convergence analysis and statistical significance testing. CAAS-enabled warehouse configurations achieve the hypothesized "better or comparable economic feasibility" through demonstrated cost parity with the 0-warehouse baseline. Meanwhile, CAAS provides environmental and operational benefits.

The sustainability performance meets hypothesis expectations with statistically significant emissions reductions and establishment of satellite refurbishment capability (65-191 satellites serviced over 30 years depending on the configuration). These environmental advantages persist robustly across all performance percentiles and uncertainty scenarios, demonstrating fundamental architectural benefits rather than scenario-dependent outcomes. The mechanisms hypothesized to deliver these benefits receive validation.

Organizations facing emissions constraints, environmental regulations, or corporate sustainability commitments will find compelling justification in the demonstrated significant reductions across emissions and refurbishment metrics. The achievement of cost parity means these environmental benefits come without economic penalties over the long run, enabling sustainability leadership without substantial financial sacrifice.

Overall, the hypothesis is substantiated: multi-echelon collection warehouses with strategic satellite collection and ADR deployment, collaborative maneuvering through RPO-capable satellites, and Earth-return operations using reusable launch vehicles provide comparable economic feasibility while delivering significant sustainability improvements. This establishes CAAS as a compelling solution for organizations operating under environmental constraints or prioritizing long-term sustainability, and positions the architecture as economically viable even for organizations with less immediate environmental pressures. The achievement of cost comparability combined with sustainability superiority represents a threshold accomplishment for real-world implementation of circular space systems and validates the fundamental value proposition of warehouse-based satellite collection and servicing infrastructure.

4.3 H2 Experimentation: Flexibility

Experiment set 2 considers how the CAAS concept could be improved with flexible, incremental infrastructure deployment. To recap, the relevant research questions and hypothesis are as follows:

Research Question 3: Flexibility Options

Which flexible option, or set of flexible options, provides the greatest economic value and environmental benefit for OOS providers under uncertainty, compared to inflexible infrastructures?

Hypothesis 2

If the flexibility framework models CAAS system evolution, captures uncertainty, and models interactive decision-making, it will identify which flexible enablers and strategies for both servicing infrastructure and satellite constellation further improve the economic feasibility and sustainability metrics of the CAAS infrastructure compared to infrastructures that only rely on the over-population sparing strategy.

The supporting condition for hypothesis 2 is that the model identifies statistically significant flexible enablers or strategies that give CAAS infrastructures an advantage over baseline configurations in terms of both economic feasibility and sustainability metrics. Therefore, the model includes flexible strategies triggered by decision rules, as highlighted in Figure 4.7.

Experiment 2 offers equivalence to the real-world problem by modeling the capital investment decisions that infrastructure operators face when deploying OOS systems under uncertainty. Rather than assuming all infrastructure capabilities are available from day one, this experiment reflects how real infrastructure projects are often deployed: incrementally, based on demonstrated return on investment and evolving




Figure 4.7: Experiment 2 Methodology to Support Hypothesis 2

market conditions. When flexibly deployed, each technology option is triggered by decision rules that compare the anticipated benefits of an investment with its fixed and recurring costs, therefore mirroring how operators would evaluate whether to expand capabilities (like adding warehouse capability, deploying RPO-enabled satellites) based on observed constellation needs and cost trajectories. Just like in Experiment 1, Experiment 2 features various future scenarios to show how these options perform over a variety of futures and which provide consistent, measurable benefits. This experiment addresses gaps from literature by considering the value of flexibility in the rollout of the OOS infrastructure itself, not just the flexible value that OOS provides to the satellite constellation operator. While prior work has examined how on-orbit servicing benefits satellite operators, no studies have systematically evaluated which infrastructure elements should be deployed initially versus conditionally based on evolving cost conditions. This experiment determines which technologies over-extend the infrastructure versus which enable its future productivity, revealing whether technologies such as RPO-capable satellites are best deployed from the very start or deployed later in the simulation under favorable conditions. Additionally, Experiment 2 includes sensitivity studies on uncertain variables as well as aspects of the CONOPS, such as the ability to deploy ADR vehicles in space, to provide further insight into the conditions that make the CAAS concept attractive. While ideally all uncertain variables and estimated values would receive sensitivity testing, for the purpose of demonstrating the framework's capability to conduct sensitivity tests while

maintaining thesis scope, this experiment focuses on a selection of parameters that are anticipated to be top influencers on the results. Overall, Experiment 2 seeks to quantify the conditions and flexible strategies/enablers that make CAAS attractive from both the financial and environmental perspectives, specifically identifying optimal deployment timing for each infrastructure element.

4.3.1 Configuration Set

Compared to the previous experiment set, the satellite upgrade (Rf - refuelable), warehouse upgrade (Upg - upgraded), and RPO configuration variables include the possibility for dynamic deployment of the option (config = 1), along with the option never being available (config=0) and the option being immediately available (config=2), as shown below in Table 4.7. Table 4.1 explains the dependencies of each configuration while Table 4.8 illustrates the combinations that form the configuration set for this set of experiments. Since Experiment 1 uncovered that 1 warehouse is the cost-superior CAAS configuration, the flexibility experiments focus on one initial warehouse.

Table 4.7: Key Changes from Previous Configuration Logic

Parameter	Previous Logic	Updated Logic
<code>depot_config</code>	{0,1,2,3}	{0,1}
<code>refuelable_config</code>	{0,2} only	{0,1,2}
<code>upgraded_config</code>	{0,2} when refuelable = 2, else 0	{0,1,2} when refuelable $\in \{1,2\}$, else 0
<code>rpo_config</code>	{0,2} only	{0,1,2}
<code>add_new_wh_config</code>	{0} only	{0,1}

The analysis indicates that 80 uncertainty seeds are adequate for providing statistically significant results across all primary outputs (cost, emissions, number of refurbishments).

Table 4.8: Configuration Combination Summary

Configuration Type	Count	Key Characteristics
Baseline (No Depot)	1	<code>depot_config = 0</code> <code>CAAS_config = 0</code> All CAAS-related params = 0
CAAS without Sat Upgrades	6	<code>depot_config = 1</code> <code>CAAS_config = 1</code> <code>refuelable_config = 0</code> <code>upgraded_config = 0</code> <code>add_new_wh_config ∈ {0,1}</code> <code>rpo_config ∈ {0,1,2}</code> (6 combinations: 2 warehouse × 3 rpo)
CAAS with Optional Sat Upgrades	18	<code>refuelable_config = 1</code> <code>upgraded_config ∈ {0,1,2}</code> <code>add_new_wh_config ∈ {0,1}</code> <code>rpo_config ∈ {0,1,2}</code> (18 combinations: 3 × 2 × 3)
CAAS with Immediately Available Sat Upgrades	18	<code>refuelable_config = 2</code> <code>upgraded_config ∈ {0,1,2}</code> <code>add_new_wh_config ∈ {0,1}</code> <code>rpo_config ∈ {0,1,2}</code> (18 combinations: 3 × 2 × 3)
Total Combinations	43	Per scenario iteration

4.3.2 VARG Results for Experiment 2

This section presents comprehensive results of the Value at Risk and Gain (VARG) analysis combined with statistical significance testing and ranking convergence assessment for multiple performance metrics across different configurations with flexible infrastructure update options. The analysis evaluated 43 unique configurations across 80 scenarios, tabulated in Table 4.9, across seven key performance metrics, incorporating additional infrastructure flexibility parameters including warehouse additions, refuelable satellite options, and dynamic infrastructure scaling. The framework enables analysis of superior commitment strategies across different services, revealing when pre-commitment outperforms flexible deployment or vice versa.

Table 4.9 presents the settings for the top-performing configurations based on ranking convergence analysis across 80 uncertain scenarios. In all scenarios, the option to add a new warehouse is never exercised, so `add_new_wh_config` is excluded from the results. Ranking convergence analysis indicates that for average total cost, the rank order is stable and converged for nearly all configurations, with the 0-warehouse baseline maintaining rank 1 and the flexible CAAS configuration (D1, Rf1, Up0, RP2) achieving rank 2.

The introduction of flexible options reveals important insights through ranking convergence analysis rather than through statistically significant performance differences, as illustrated by the small differences in the VARG plot, provided in Figure 4.8. The cost-superior CAAS configuration is (D1, C1, Rf1, Up0, RP2, P0), which represents a 0.43% cost increase compared to the 0-warehouse baseline, while the cost-superior inflexible CAAS configuration (D1, C1, Rf0, Up0, RP2, P0) shows a 0.67% cost increase relative to baseline. However, these differences are not statistically significant given the 80 uncertainty scenarios analyzed; establishing measurable cost differences for such a small percent would require on the order of 1,000 scenarios.

The rank convergence analysis provides confidence that flexible satellite deploy-

Table 4.9: Configuration Parameter Settings for All Depots - Flexible

Configuration	Depot	Upgraded	Refuelable	RPO	CAAS
0-Warehouse Baseline	0	0	0	0	0
Best Cost No-Flexibility	1	0	0	2	1
Best for Total Cost & Op Costs & Weighted Multi-Obj	1	0	1	2	1
Best for Total Emissions	1	2	1	2	1
Best for Cost/Refurbish & Sats Refurb. In Space	1	2	2	2	1
Best for Total Num Sat Refurb	1	1	1	2	1

ment ($Rf=1$) paired with pre-initialized RPO-capable satellites represents the top-performing CAAS configuration, though the absolute magnitude of improvement cannot be precisely quantified with the current sample size. This finding highlights a critical distinction: while ranking stability indicates which configurations perform best relative to each other, determining the exact performance gap requires substantially more scenarios.

Experimental Observation 2.1

Ranking convergence analysis indicates that flexible satellite upgrade deployment ($Rf=1$) paired with pre-initialized RPO-capable satellites consistently ranks as the top-performing CAAS configurations for various metrics across uncertainty scenarios. However, with 80 scenarios, the absolute performance differences (such as 0.427% vs 0.67% avg cost increase) are too small to be measured with confidence.

For average emissions, the ranking of the top-performing configuration (D1, C1, Rf1, Up2, RP2, P0) is stable, though rankings at the 5th and 95th percentiles do not converge, indicating continued sensitivity to extreme outcomes. This configuration

achieves an average of 63 refurbished satellites in total and 19 satellites refurbished in space. Notably, all top-performing configurations share RPO=2, indicating the critical importance of pre-initialized RPO capability for improving refurbishment rates, cost, and emissions.




Figure 4.8: Total Cost for Flexible CAAS vs. Inflexible CAAS vs. Baseline for Experiment 2: Ranking convergence indicates stable relative ordering despite lack of statistical significance in absolute differences

Like the cost results, improvements in emissions associated with flexibility are not measurable with 80 scenarios. The ranking convergence analysis provides confidence in the relative ordering of configurations based on average performance but cannot precisely quantify the magnitude of emissions reductions. This analysis further confirms that the number of depots and pre-initialized RPO-capability emerge as the prominent factor affecting both emissions and refurbishment metrics, with incremental deployment playing a secondary role compared to infrastructure scale.

Table 4.10: Experiment 2: Total Cost

Top Performing Metric	Configuration	Percentile	Total Cost Value (\$)	Total Cost CI Lower	Total Cost CI Upper	Total Cost MoE%	Total Cost % vs Baseline	Total Cost p-value (Baseline)	Total Cost Sig? (Baseline)	Total Cost % vs No-CAAS	Total Cost p-value (No-CAAS)	Total Cost Sig? (No-CAAS)
0-Warehouse Baseline	d:0 up:0 rf:0 RPO:0 CAAS:0	avg	4.43E+09	4.34E+09	4.52E+09	2.090	N/A	N/A	N/A	-0.668	0.632	No
	5th	3.97E+09	3.89E+09	4.03E+09	1.754	N/A	N/A	N/A	-1.896	0.632	No	
	95th	5.25E+09	5.02E+09	5.69E+09	6.375	N/A	N/A	N/A	1.684	0.632	No	
Best Cost No-Flexibility	d:1 up:0 rf:0 RPO:2 CAAS:0	avg	4.46E+09	4.38E+09	4.54E+09	1.832	0.673	0.632	No	N/A	N/A	N/A
	5th	4.05E+09	4.03E+09	4.10E+09	0.831	1.933	0.632	No	N/A	N/A	N/A	
	95th	5.17E+09	4.99E+09	5.56E+09	5.516	-1.656	0.632	No	N/A	N/A	N/A	
Best for Total Cost, Op Cost, Weighted Multi-Objective	d:1 up:2 rf:1 RPO:2 CAAS:1	avg	4.45E+09	4.37E+09	4.53E+09	1.852	0.427	0.762	No	-0.244	0.852	No
	5th	4.06E+09	4.02E+09	4.09E+09	0.851	2.148	0.762	No	0.211	0.852	No	
	95th	5.17E+09	4.93E+09	5.55E+09	6.024	-1.658	0.762	No	-0.002	0.852	No	
Best for Total Emissions	d:1 up:2 rf:1 RPO:2 CAAS:1	avg	4.48E+09	4.40E+09	4.56E+09	1.847	1.158	0.412	No	0.482	0.713	No
	5th	4.08E+09	4.05E+09	4.11E+09	0.820	2.760	0.412	No	0.812	0.713	No	
	95th	5.18E+09	4.96E+09	5.57E+09	5.900	-1.341	0.412	No	0.320	0.713	No	
Best for Cost/Refurbish, Sats Refurb. In Space	d:1 up:2 rf:2 RPO:2 CAAS:1	avg	4.53E+09	4.44E+09	4.61E+09	1.864	2.150	0.132	No	1.467	0.269	No
	5th	4.12E+09	4.09E+09	4.14E+09	0.666	3.663	0.132	No	1.697	0.269	No	
	95th	5.27E+09	4.98E+09	5.70E+09	6.850	0.262	0.132	No	1.950	0.269	No	
Best for Total Num Sat Refurb	d:1 up:1 rf:1 RPO:2 CAAS:1	avg	4.45E+09	4.37E+09	4.54E+09	1.868	0.545	0.700	No	-0.127	0.923	No
	5th	4.06E+09	4.02E+09	4.10E+09	1.033	2.252	0.700	No	0.314	0.923	No	
	95th	5.19E+09	4.94E+09	5.55E+09	5.861	-1.155	0.700	No	0.509	0.923	No	

Table 4.11: Experiment 2: Total NOx Emissions

Top Performing Metric	Configuration	Percentile	Total Emissions Value (kg)	Total Emissions CI Lower	Total Emissions CI Upper	Total Emissions MoE%	Total Emissions % vs Baseline	Total Emissions p-value (Baseline)	Total Emissions Sig? (Baseline)	Total Emissions % vs No-Flexibility	Total Emissions p-value (No-Flexibility)	Total Emissions Sig? (No-Flexibility)
0-Warehouse Baseline	d0 up:0 rf:0 RPO:0 CAAS:0	avg	3.49E+06	3.41E+06	3.56E+06	2.071	N/A	N/A	N/A	6.861	0.000	Yes
		5th	2.98E+06	2.83E+06	3.04E+06	3.527	N/A	N/A	N/A	12.874	0.000	Yes
		95th	3.95E+06	3.91E+06	4.01E+06	1.191	N/A	N/A	N/A	5.531	0.000	Yes
Best Cost No-Flexibility	d:1 up:0 rf:0 RPO:2 CAAS:0	avg	3.26E+06	3.19E+06	3.33E+06	2.255	-6.421	0.000	Yes	N/A	N/A	N/A
		5th	2.64E+06	2.60E+06	2.86E+06	5.019	-11.406	0.000	Yes	N/A	N/A	N/A
		95th	3.74E+06	3.68E+06	3.77E+06	1.150	-5.241	0.000	Yes	N/A	N/A	N/A
Best for Total Cost, Op Cost, Weighted Multi-Objective	d:1 up:2 rf:1 RPO:2 CAAS:1	avg	3.27E+06	3.18E+06	3.35E+06	2.588	-6.237	0.000	Yes	0.196	0.910	No
		5th	2.62E+06	2.46E+06	2.79E+06	6.283	-12.138	0.000	Yes	-0.826	0.910	No
		95th	3.87E+06	3.74E+06	3.92E+06	2.282	-1.981	0.000	Yes	3.441	0.910	No
Best for Total Emissions	d:1 up:2 rf:1 RPO:2 CAAS:1	avg	3.25E+06	3.17E+06	3.33E+06	2.427	-6.645	0.000	Yes	-0.239	0.886	No
		5th	2.55E+06	2.48E+06	2.74E+06	5.038	-14.332	0.000	Yes	-3.303	0.886	No
		95th	3.74E+06	3.68E+06	3.86E+06	2.445	-5.247	0.000	Yes	-0.006	0.886	No
Best for Cost/Refurbish, Sats Refurb. In Space	d:1 up:2 rf:2 RPO:2 CAAS:1	avg	3.27E+06	3.19E+06	3.34E+06	2.356	-6.272	0.000	Yes	0.160	0.923	No
		5th	2.66E+06	2.54E+06	2.84E+06	5.651	-10.788	0.000	Yes	0.697	0.923	No
		95th	3.80E+06	3.68E+06	3.82E+06	1.766	-3.663	0.000	Yes	1.665	0.923	No
Best for Total Num Sat Refurb	d:1 up:1 rf:1 RPO:2 CAAS:1	avg	3.26E+06	3.18E+06	3.35E+06	2.513	-6.346	0.000	Yes	0.080	0.962	No
		5th	2.55E+06	2.47E+06	2.80E+06	6.301	-14.445	0.000	Yes	-3.431	0.962	No
		95th	3.85E+06	3.68E+06	3.89E+06	2.702	-2.351	0.000	Yes	3.050	0.962	No

Table 4.12: Experiment 2: Number of Refurbishments

Top Performing Metric	Configuration	Percentile	Num Sat Refurb Value	Num Sat Refurb CI Lower	Num Sat Refurb CI Upper	Num Sat Refurb MoE%	Num Sat Refurb % vs No-Flexibility	Num Sat Refurb p-value (No-Flexibility)	Num Sat Refurb Sig? (No-Flexibility)
0-Warehouse Baseline	d:0 up:0 rf:0 RPO:0 CAAS:0	avg	0.00	0.00	0.00	N/A	N/A	N/A	N/A
		5th	0.00	0.00	0.00	N/A	N/A	N/A	N/A
		95th	0.00	0.00	0.00	N/A	N/A	N/A	N/A
Best Cost No-Flexibility	d:1 up:0 rf:0 RPO:2 CAAS:0	avg	64.89	53.60	76.17	17.389	N/A	N/A	N/A
		5th	0.00	0.00	0.00	0.000	N/A	N/A	N/A
		95th	143.30	123.00	192.00	24.075	N/A	N/A	N/A
Best for Total Cost, Op Cost, Weighted Multi-Objective	d:1 up:2 rf:1 RPO:2 CAAS:1	avg	66.20	54.59	77.81	17.545	2.023	0.872	No
		5th	0.00	0.00	0.00	0.000	-	-	N/A
		95th	142.95	125.00	165.00	13.991	-0.244	0.872	No
Best for Total Emissions	d:1 up:2 rf:1 RPO:2 CAAS:1	avg	62.55	53.67	71.43	14.194	-3.602	0.746	No
		5th	0.00	0.00	8.55	0.000	-	-	N/A
		95th	128.00	118.00	143.00	9.766	-10.677	0.746	No
Best for Cost/Refurbish, Sats Refurb. In Space	d:1 up:2 rf:2 RPO:2 CAAS:1	avg	69.34	60.79	77.89	12.334	6.858	0.533	No
		5th	3.00	1.00	30.00	483.333	inf	0.533	No
		95th	122.30	117.00	146.00	11.856	-14.655	0.533	No
Best for Total Num Sat Refurb	d:1 up:1 rf:1 RPO:2 CAAS:1	avg	69.71	59.45	79.97	14.720	7.436	0.530	No
		5th	0.00	0.00	8.75	0.000	-	-	N/A
		95th	136.35	128.05	163.00	12.816	-4.850	0.530	No

Experiment 2 Ranking Convergence and Statistical Significance

The statistical analysis of Experiment 2 reveals a fundamental limitation: differences in total cost and emissions between flexible configurations, the baseline, and top no-flexibility configurations are not statistically significant at any percentile (5th, average, or 95th) given the 80 uncertainty scenarios analyzed. While flexible deployment ($R_f=1$) demonstrates consistent improvement in rank order, these improvements are not measurable with the current sample size. Establishing statistical significance for the observed performance differences would require on the order of 1,000 scenarios, which would require cluster computing since each simulation takes 1-1.5 minutes to run.

This finding necessitates reliance on ranking convergence analysis rather than traditional hypothesis testing. The convergence framework provides confidence in the relative ordering of configurations while acknowledging that absolute performance differences remain too small to quantify precisely. For total cost, rankings converged for nearly all configurations at the average, with the 0-warehouse baseline maintaining rank 1 and (D1, R_f1 , Up0, RP2) achieving stable rank 2. At the 5th percentile, the top 3 configurations show converged rankings, though 95th percentile rankings have not yet stabilized after 80 scenarios, indicating continued sensitivity to worst-case outcomes.

For emissions, the rank-1 configuration (D1, R_f1 , Up2, RP2) has stabilized at the average, though neither the 5th nor 95th percentile rankings converge after 80 scenarios. Again, all top-performing configurations share $RPO=2$, continuing the pattern from Experiment 1 even though this experiment included the option for flexible RPO -capability deployment.

The rank order provides confidence that flexible satellite upgrades paired with pre-initialized RPO -capable satellites represent the superior CAAS configuration for cost management, but determining whether savings this translates to requires substantially

more scenarios. This distinction is critical for decision-making: ranking convergence guides technology selection and policy priorities, while statistical significance would enable precise cost-benefit analysis and investment justification.

In addition to analyzing rank order convergence and hypothesis testing on configuration performance comparison, this experiment also uses regression-based variance testing to determine the effect of each individual configuration variable. Statistical significance is assessed at $p < 0.05$ with polynomial features up to degree 3. Starting off with total cost, the satellite serviceability upgrade variable (Rf) and warehouse upgrade variable (Upg) are significant for all percentiles. Meanwhile, the RPO configuration is not statistically significant for the 5th total cost. This suggests that RPO primarily affects average performance and tail risk rather than extend opportunity

For total emissions, the satellite serviceability upgrade flexibility variable (Rf) demonstrates a nuanced relationship with performance. While not statistically significant for average or 95th percentile emissions, the satellite upgrade variable (Rf) becomes significant at the 5th percentile, implying the satellite serviceability makes an impact on extending the opportunity to reduce emissions. The RPO configuration is significant for average and 5th percentile emissions, but not the 95th percentile. Lastly, warehouse upgrades (Upg) are not statistically significant at any percentile for total emissions.

Total refurbishments rely on all configuration variables, with the exception of the 95th percentile, which solely depends on the number of warehouses and RPO-capability. Meanwhile, total cost divided by total refurbishments depends on all configuration variables except for the 5th percentile, which only replies on the number of depots. Overall, approximating variable influence using this regression-based variance analysis aligns with ranking stability analysis; for instance, emissions ranking becomes more stable when the warehouse upgrade variable is excluded from the analysis.

The convergence analysis yields several key insights that shape interpretation of the results. First, the ranking framework successfully identifies relative performance ordering even when absolute differences are too small for statistical significance. This enables technology selection and policy prioritization based on converged rankings while acknowledging uncertainty in precise performance quantification. Second, the model reveals that incremental deployment flexibility is not as significant a driver in reducing emissions as the fundamental choice of depot count, highlighting where policy attention should focus. Third, Earth-based refurbishment emerges as the critical factor for the NOx emissions metric since adding space-based refurbishment has marginal impact.

The frequent lack of statistical significance, despite stable rankings, has important implications for technology adoption strategies. Flexibility investments may appear unjustified under traditional cost-benefit analysis requiring precise performance quantification, yet the ranking convergence provides strong evidence for their relative superiority. This suggests that decision-making frameworks should incorporate ranking-based evidence when sample sizes limit statistical power, particularly for complex system-of-systems analyses where the cost of generating thousands of scenarios may be prohibitive. Future work should focus on increasing scenario counts for configurations near the top of each ranking to achieve statistical significance for precise performance differences, while the current analysis provides sufficient confidence for initial technology guidance and policy direction.

4.3.3 Sensitivity Testing

Uncertain Variable Sensitivity Test

To assess the robustness of the results to key sources of uncertainty, Experiment 2 was re-run under various alternative conditions, each isolating a different uncertain variable or operational condition. To manage computational expense while providing

a sufficient sample size for stable ranking, these tests use 20 scenarios each. While 20 scenarios are sufficient to establish stable cost rankings between configurations, they are insufficient to quantify absolute performance differences between sensitivity test results and nominal cases. Therefore, the cost and emissions differences reported relative to the 0-warehouse baseline (sensitivity test vs. nominal) provide notional insights into the magnitude of effects rather than statistically significant measurements. Ideally, all assumptions should be perturbed to better understand their impact on CAAS performance and each test should use 1000s of scenarios to quantify absolute differences, but in the interest of scope, this thesis focuses on rank order comparison for key operational uncertainties, like failure rates, as well as key cost uncertainties, like ADR and servicing costs. Table 4.13 shows which configuration variables were varied for these tests. Table 4.14, Table 4.15, and Table 4.16 summarize the results.

Table 4.13: Sensitivity Test Configuration Parameters

Parameter	Values	Dependency	Description
Number of warehouses	{0,1}	None	Depot infrastructure configuration: 0 = No warehouse (baseline); 1 = Number of warehouses.
CAAS Mode	{0,1}	# warehouses \neq 0	Collection-as-a-Service availability: 0 = Disabled; 1 = Enabled.
Refuelable/Repairable Upgrades	{0,1,2}	# warehouses \neq 0 and CAAS = 1	0 = No upgrades; 1 = flexible upgrades; 2 = Immediate upgrades.
Warehouse Upgrades	{0,1,2}	Upgraded satellites = 2	Warehouse upgrade strategy: 0 = No upgrades; 1 = flexible upgrades; 2 = Immediate upgrades.
RPO-Capable Satellites	{0,2}	# warehouses \neq 0 and CAAS = 1	Rendezvous and Proximity Operations capability: 0 = Disabled; 2 = Available.

Table 4.14: Sensitivity analysis results for key uncertain variables (top) and failure modeling, operations, and management (bottom)

Sensitivity Test	Top CAAS Cost Config	Top Emissions Config	Cost Diff. vs. Nominal	Emis. Diff. vs. Nominal
Technology Obsolescence Removed	(D:1, Rf:1, Upg:0, RPO:2)	(D:1, Rf:1, Upg:1, RPO:2)	+0.6%, 0.6%	-6.8%, -5.3%
Launch Cost Held Constant	(D:1, Rf:1, Upg:0, RPO:2)	(D:1, Rf:1, RPO:2)	+1.0%, 0.6%	-5.9%, -5.3%
Satellite Cost Held Constant	(D:1, Rf:1, Upg:0, RPO:2)	(D:1, Rf:1, Upg:0, RPO:2)	-0.6%, 0.6%	-7.9%, -5.3%
0.75X initial failure rate	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	+1.8%, 0.6%	-7.5%, -7.3%
0.5X initial failure rate	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 1, RPO: 2, P: 0)	+3.0%, 0.6%	-6.4%, -7.3%
ADR launched exclusively from Earth	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	+1.4%, 0.8%	-8.1%, -6.4%
No Reusable Second Stage Vehicles	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D:1, Rf:2, Upg:2, RPO:2, P:0)	+0.8%, 0.5%	+0.5%, -3.6%

247

Note: Sensitivity Cost and Emission difference columns show the difference between the top cost CAAS config compared to the 0-warehouse baseline, comparing the sensitivity result (left) to a nominal case without the sensitivity test applied (right). Bold configuration tags indicate that CAAS configuration with sensitivity test applied outperforms the 0-warehouse baseline in terms of cost (stable 1st place rank). Configuration tags include all configuration variables that could converge to a stable ranking within 20 scenarios.

The top half of Table 4.14 includes sensitivity tests on key uncertain variables. The removal of technology obsolescence yielded no difference in configuration ranking compared to the original configuration, as the simulation algorithm inherently prevents obsolete satellites from remaining active for extended periods, which means the revenue loss from obsolete satellites has a minimal impact on total cost. When launch costs were held at their initial values, the superior configuration remained unchanged, demonstrating 2nd place ranking with cost differences staying within 1% of the baseline. Notably, when satellite costs were held constant, the top CAAS cost configuration achieved stable 1st place ranking, outperforming the 0-warehouse baseline by -0.6%. This suggests that the CAAS concept becomes particularly attractive if satellite costs decrease less over time than anticipated in this framework, as the relative value of refurbishment infrastructure increases when new satellites remain expensive.

The bottom half of Table 4.14 includes sensitivity tests on failure modeling, operations, and failure management. All four tests maintained stable 2nd place cost ranking, with the CAAS concept maintaining its measurable emissions advantage over the baseline. The first two tests vary failure rate modeling by reducing the assumed initial failure rate (1 satellite/year) by 3/4 and 1/2. Initial failure rate assumptions demonstrated the most pronounced effects on cost competitiveness among the uncertain variable tests, though the top CAAS configuration remained stable. The standard model in this thesis assumes 1 satellite failure per year at the start of the simulation. Reducing the initial failure rate to 0.75 failures per year increased the cost difference to +1.8%, while further reduction to 0.5 failures per year yielded +3.0%, compared to +0.6% under the baseline failure rate. This sensitivity reveals a critical insight: the CAAS concept becomes significantly less cost- competitive when satellite reliability improves, as fewer failures diminish the value of the refurbishment infrastructure. This relationship highlights that the economic case for CAAS is strongest

in environments with higher satellite failure rates. The third test removes the ability to deploy ADR vehicles from the warehouses and restricts all ADR missions to launch directly from Earth. Launching ADR entirely from Earth rather than including depot-based deployment increased the cost difference to +1.4% (versus +0.8% with nominal, depot-based ADR), highlighting the operational cost savings enabled by depot-based ADR deployment. Interestingly, when ADR vehicles are launched exclusively from Earth, there are slightly fewer emissions (-8.1%) than the nominal CAAS case (-6.4%). While not measurable given the number of cases, this shows how warehouse-based deployment of ADR vehicles may provide more cost benefits than emissions benefits: reusing ADR vehicles saves money but requires launches to sustain ADR-refueling operations in space. The 4th test removes all reusable second stage vehicles from all configurations, which causes a slight increase in cost difference relative to the nominal case and a slight *increase* in emissions relative to the 0-warehouse baseline (which also includes the no-reusable second stage vehicle effect).

Importantly, across all uncertain variable sensitivity tests, the cost remained within 3% of the 0-warehouse baseline, and with the exception of the no-reusable second stage test, no test eliminated the emissions benefits of CAAS. Among all uncertain variables tested, failure rate assumptions emerged as a strong driver of cost performance, merited further attention in future work. The stability of configuration rankings and the absence of any extreme changes in cost or emissions reinforce confidence that the CAAS concept remains economically viable and environmentally beneficial across a wide range of uncertain conditions.

Deterministic Variable Sensitivity Test

To assess the robustness of the results to key deterministic assumptions, Experiment 2 was re-run with systematic perturbations to operational cost parameters using 20 uncertain scenarios just as before. Table 4.15 and Table 4.16 present the results of

these sensitivity tests, varying cost parameters with various multipliers. Two specific conditions led to the CAAS configuration achieving stable 1st place cost ranking (outperforming the 0-warehouse baseline):

- ADR Cost Min Fraction reductions (0.5X and 0.25X): When the minimum achievable fraction of present-day ADR vehicle cost through learning curve effects is reduced to 50% or 25% of baseline assumptions, the CAAS concept achieves cost superiority over the baseline (-1.0% and -2.0%, respectively). This indicates that continued improvements in ADR vehicle manufacturing and deployment efficiency could make CAAS economically superior.
- Earth-based Refuel and Repair Streamlining (0.5X and 0.25X multipliers): When satellites are made to be refuelable or repairable, Earth-based refurbishment costs decrease by some uncertain factor. When these factors are reduced by 50% or 25% (indicating more efficient Earth-based refurbishment processes), the CAAS concept achieves stable 1st place ranking. This suggests that innovations in Earth-based satellite servicing design could tip the economic balance in favor of CAAS.

All other cost perturbations resulted in stable 2nd place cost ranking for the CAAS configuration, with cost differences remaining within +0.5% to +0.9% of the 0-warehouse baseline. These small margins indicate near-competitive performance across a wide range of operational cost assumptions. Critically, even the most pessimistic cost perturbations (e.g., 1.5X increases in refuel, refurbish, or repair costs) did not cause CAAS costs to "blow up" or diverge significantly from baseline, and all scenarios maintained sufficient emissions reductions.

The superior CAAS configuration across all perturbation scenarios remained consistent: one depot (D1), CAAS (C1), flexible satellite upgrades (Rf1), no warehouse upgrades (Up0), and RPO-capable satellites (RP2). The consistent emergence of this

configuration, combined with the absence of catastrophic cost modes under pessimistic assumptions, reinforces confidence that the recommended architecture is robust to parameter uncertainty within reasonable bounds.

Experimental Observation 2.2

The top ranked CAAS configuration (D:1, Rf:1, Upg:0, RPO:2) stays constant through all sensitivity tests. While most scenarios show stable 2nd place cost ranking with small differences from the 0-warehouse baseline, three conditions achieve stable 1st place ranking: static satellite costs, reduced ADR minimum costs (learning curve improvements), and streamlined refuel/repair Earth-based processes. Failure rate emerges as the strongest driver of cost performance among uncertain variables, with optimistic reliability assumptions degrading CAAS cost-competitiveness. Importantly, no sensitivity test caused extreme costs or eliminated emissions benefits (aside from removing reusable second stage vehicles), reinforcing confidence in the CAAS concept's robustness.

Risk Management Implications These sensitivity results provide important guidance for risk-conscious decision-making. The stability of depot number, RPO, and satellite upgrade configurations across all perturbation scenarios suggests these represent high-confidence decisions that should remain consistent regardless of moderate parameter uncertainty.

Failure rate emerges as the most critical uncertain variable for CAAS cost-competitiveness. The 2.4 percentage point cost swing between baseline and optimistic reliability scenarios indicates that stakeholders should carefully assess expected satellite failure rates in their specific operational environment before committing to CAAS infrastructure. The economic case for CAAS is strongest when initial satellite reliability starts on the order of 1 failure per year, as higher failure rates maximize the failure-collection benefits of the CAAS concept.

Table 4.15: Deterministic sensitivity analysis results for operational cost parameters (Part 1)

Sensitivity Test	Top CAAS Cost Config	Top Emissions Config	Cost Diff. vs. Nominal	Emis. Diff. vs. Nominal
ADR cost Min Fraction 0.5X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	-1.0%, 0.7%	-7.5%, -5.2%
ADR cost Min Fraction 0.25X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Upg: 1, RPO: 2, P: 0)	-2.0%, 0.7%	-6.8%, -5.2%
Earth-based Refuel and Repair Stream. $\times 0.5$	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, RPO: 2, P: 0)	+0.5%, 0.7%	-5.9%, -5.2%
Earth-based Refuel and Repair Stream. $\times 0.25$	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, RPO: 2, P: 0)	+0.4%, 0.7%	-7.0%, -5.2%
ADR Launch Cost 1.25X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 0, Upg: 0, RPO: 2, P: 0)	+0.9%, 0.6%	-6.5%, -6.4%
ADR Launch Cost 1.50X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, RPO: 2, P: 0)	+0.9%, 0.6%	-5.2%, -6.4%
Refuel Cost 1.25X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 2, Upg: 1, RPO: 2, P: 0)	+0.6%, 0.5%	-7.4%, -7.9%
Refuel Cost 1.50X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, RPO: 2, P: 0)	+0.6%, 0.5%	-8.3%, -7.9%
Refurbish Cost 1.25X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 2, Upg: 0, RPO: 2, P: 0)	+0.8%, 0.7%	-5.7%, -5.6%

Note: Sensitivity Cost and Emission difference columns show the difference between the top cost CAAS config compared to the 0-warehouse baseline, comparing the sensitivity result (left) to a nominal case without the sensitivity test applied (right). Bold configuration tags indicate that CAAS configuration with sensitivity test applied outperforms the 0-warehouse baseline in terms of cost (stable 1st place rank). Configuration tags include all configuration variables that could converge to a stable ranking within 20 scenarios.

Table 4.16: Deterministic sensitivity analysis results for operational cost parameters (Part 2)

Sensitivity Test	Top CAAS Cost Config	Top Emissions Config	Cost Diff. vs. Nominal	Emis. Diff. vs. Nominal
Refurbish Cost 1.50X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 0, Upg: 0, RPO: 2, P: 0)	+0.7%, 0.7%	-5.8%, -5.6%
Repair Cost 1.25X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, RPO: 2, P: 0)	+0.7%, 0.7%	-6.6%, -5.6%
Repair Cost 1.50X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 2, RPO: 2, P: 0)	+0.7%, 0.7%	-5.4%, -5.6%
ADR Op. Cost 1.25X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 2, Upg: 2, RPO: 2, P: 0)	+0.8%, 0.6%	-6.3%, -6.3%
ADR Op. Cost 1.50X	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	+0.8%, 0.6%	-7.6%, -6.3%
Return Cost 1.0 X launch	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 2, Upg: 1, RPO: 2, P: 0)	+0.8%, 0.7%	-6.6%, -7.1%
Return Cost 1.5X launch	(D: 1, Rf: 1, Upg: 0, RPO: 2, P: 0)	(D: 1, Rf: 1, Upg: 2, RPO: 2, P: 0)	+0.8%, 0.7%	-6.8%, -7.1%

Note: Sensitivity Cost and Emission difference columns show the difference between the top cost CAAS config compared to the 0-warehouse baseline, comparing the sensitivity result (left) to a nominal case without the sensitivity test applied (right). Bold configuration tags indicate that CAAS configuration with sensitivity test applied outperforms the 0-warehouse baseline in terms of cost (stable 1st place rank). Configuration tags include all configuration variables that could converge to a stable ranking within 20 scenarios.

The analysis also reveals two controllable pathways to cost superiority for CAAS: (1) improvements in ADR vehicle cost reduction through learning curve effects, and (2) innovations that streamline Earth-based satellite servicing design. For cost-focused stakeholders, prioritizing R&D investments in these areas could enhance CAAS economic competitiveness. In contrast, instability in emissions-superior configurations across sensitivity tests (related to the warehouse upgrade configuration variable) suggests that warehouse upgrades present risk for emissions-focused stakeholders, since the benefits of initialized or flexibly-deployed warehouse upgrades are not consistent.

4.3.4 Experimental Support for Hypothesis 2

The central hypothesis posits that: *If the flexibility framework models CAAS system evolution, captures uncertainty, and models interactive decision-making, it will identify which flexible enablers and strategies for both servicing infrastructure and satellite constellation further improve the economic feasibility and sustainability metrics of the CAAS infrastructure compared to infrastructures that only rely on the overpopulation sparing strategy.*

This hypothesis contains three testable components: (1) identifying beneficial flexible enablers and strategies, (2) improving economic feasibility, and (3) improving sustainability metrics, all relative to baseline strategies.

Key Findings

The flexibility experiment demonstrates that flexible satellite upgrade capability ($Rf=1$) provides marginal operational value. Combined with initialized RPO capabilities ($RPO=2$), the cost-optimal flexible configuration ($D1, C1, Rf1, Up0, RP2, P0$) shows a 0.427% cost increase relative to baseline, compared to 0.67% for the best inflexible CAAS configuration ($D1, C1, Rf0, Up0, RP2, P0$). This represents an improvement of

0.243 percentage points, though these differences are not statistically significant with 80 uncertainty scenarios. The analysis relies on ranking convergence to establish relative performance ordering rather than precise quantification of benefits. Meanwhile, presence of the warehouse upgrade variable (either 1 or 2) does not have statistically significant impacts on emissions or total refurbishments. The framework reveals that incremental deployment flexibility plays a secondary role compared to infrastructure scale decisions and initial flexibility-enablers, like having RPO-capable satellites from the start. Earth-return proves critical for NOx emissions reduction, while flexible satellite upgrades provide consistent albeit small benefits across metrics.

The recommended configuration, 1 depot, flexible satellite upgrades ($Rf=1$), no warehouse upgrades ($Upg=0$), and initialized RPO ($RPO=2$), proves to be robust across the sensitivity analyses. Under deterministic cost and uncertain variable perturbations, the configuration maintains cost differences within a small percentage of the 0-warehouse baseline, with exceptions occurring under reduced failure rate conditions (+1.8% at $0.75 \times$ initial rate, +3% at $0.5 \times$ initial rate), highlighting that CAAS economic competitiveness depends critically on satellite failure frequency.

Framework Capability Assessment

The framework successfully identified and differentiated flexible enablers through VARG ranking convergence methodology combined with regression-based variance testing:

Flexible satellite serviceability ($Rf=1$): Ranking convergence indicates this provides the most consistent improvement for several metrics across scenarios. Regression-based analysis of variance shows Rf is statistically significant ($p < 0.05$) for total cost at all percentiles and for emissions at the 5th percentile, suggesting that it extends opportunity to reduce emissions in best-case scenarios.

RPO capability ($RPO=2$): Pre-initialized RPO capability emerges as criti-

cal across all top-performing configurations. Regression-based analysis of variance indicates significance for average and 5th percentile emissions, but not 95th percentile, suggesting that RPO affects average performance and downside risk rather than worst-case scenarios. All top-ranked configurations share RPO=2.

Warehouse upgrades (Upg): Regression-based analysis of variance shows warehouse upgrades are statistically significant for cost at all percentiles but not significant for emissions at any percentile. While upgraded warehouses improve space-based refurbishment throughput, space-based refurbishment does not appear to be as attractive as collection for Earth-return. Rankings become more stable when warehouse upgrade variables are excluded, suggesting warehouse upgrades only have minor importance.

Depot count: Consistently emerges as the dominant factor for emissions and refurbishment metrics. In this experiment, the simulated decision-maker never exercised the option to add warehouses given the up-front costs, confirming the need for policy intervention (addressed in Experiment 3) to achieve greater infrastructure scale and associated environmental benefits.

Top Performing Aspects of the CAAS Concept: The sensitivity test supports the finding that reusable second stage vehicles are a critical enabler for the CAAS concept, providing much of the cost and emissions savings. Additionally, the sensitivity tests indicate that ADR vehicle deployment from CAAS warehouses further improves the cost feasibility of the CAAS concept.

Economic Feasibility and Value Proposition

Flexibility demonstrates improved ranking performance relative to inflexible CAAS configurations, though the absolute magnitude of improvement (0.243 percentage points) cannot be precisely quantified with statistical confidence given the 80-scenario sample size. The ranking convergence framework provides confidence in relative or-

dering while acknowledging limitations in measuring exact performance differences, which would require approximately 1,000 scenarios to establish statistical significance.

The strategic value of flexibility lies not in dramatic cost reductions but in consistent relative performance across diverse uncertainty scenarios. Notably, sensitivity analyses demonstrate that flexibility maintains its ranking advantage even under operational cost perturbations.

However, the framework reveals a critical economic constraint: CAAS competitiveness degrades substantially under improved satellite reliability conditions. At $0.5 \times$ the baseline failure rate (0.5 failures/year initial), cost disadvantage increases to +3% versus baseline, indicating that the refurbishment infrastructure value proposition depends fundamentally on satellite failure frequency. This finding has important implications for technology adoption timing: CAAS becomes less attractive as satellite reliability improves. Based on historical evidence, 1 failure per year is a reasonable starting failure rate for the OneWeb use case, but future work that considers other constellations should be mindful of their failure rate approximations.

Hypothesis 2 Substantiation

The hypothesis is partially supported with important methodological and substantive qualifications. The framework successfully identifies flexible satellite serviceability ($R_f=1$) as the superior flexible strategy through ranking convergence analysis, even when statistical significance cannot be established for absolute performance differences. The combination of rank stability, regression-based variance testing, and sensitivity tests provide sufficient confidence for technology and CONOPs guidance despite sample size limitations, providing a credible basis for a screening tool. Overall, achieving true economic parity or advantage depends on factors beyond flexibility alone.

The hypothesis is therefore partially supported: the framework identifies beneficial

flexible enablers that should be initialized from the start and determines relative ranking of flexible options, but the economic and sustainability improvements are modest and context-dependent rather than transformative. Experiment 3 investigates whether policy combinations with flexibility can overcome these limitations and create pathways toward improved system circularity and feasibility.

4.4 H3 Experimentation: Policy

The third and final set of experiments apply policy parameters to the flexible CAAS infrastructure to determine which policy set, if any, improves sustainability without excessive costs to constellation operator and government. The related research question and hypothesis are as follows.

Research Question 4: Policy Design

Which combination and calibration of policy parameters most effectively closes the business case gap for OOS in LEO with minimal impact on overall industry costs?

Hypothesis 3

If the flexibility framework models government intervention as parametric, time-dependent “rewards” and “penalties”, then there exist scenario-dependent reward/penalty schemes that best establish economically feasible OOS infrastructures that yield better sustainability metrics than laissez-faire OOS infrastructures developed with the same incremental deployment framework.

Figure 4.9 illustrates the additions to the methodology to support experimentation for hypothesis 3. Experiment 3.1 runs a wider range of policy parameters for each proposed policy type in order to reduce the design space and identify parameters that improve total cost, emissions, and/or throughput. For this initial analysis, the satellite upgrade option and warehouse upgrade option are held constant at their flexible

option setting ($Rf=1$; $Upg=1$). Next, Experiment 3.2 directly compares the down-selected parameters from Experiment 3.1 and includes all settings for the warehouse upgrade configuration variable to assess how policy might interact with warehouse upgrade flexibility ($Rf=1$; $Upg=0,1,2$). Lastly, Experiment 3.3 takes a further down-selection of policy parameters from Experiment 3.2 and considers the effect of starting with either 1, 4 or 7 warehouses, as opposed to Experiments 3.1 and 3.2 that only consider 1 initial warehouse for CAAS configurations with policy. Experiment 3.3 includes all settings for the warehouse upgrade configuration variable and includes both flexible and pre-initialized satellite upgrades ($Rf=1,2$; $Upg=0,1,2$).




Figure 4.9: Experiment 1 Methodology to Support Hypothesis 3

Experiment 3 provides equivalence to the real-world problem by building upon the previous two experiments but now including the effect of policy implementation on the cost and benefits assessment of the CAAS concept. These experiments incorporate the various cost and reward parameters of each tested policy—all of which are based on existing or proposed policies—into the very decision-making of the infrastructure. This resembles the real world because it models how satellite operators would respond to various policy mechanisms, providing a testbed to show how proposed policies interact with the CAAS concept. Additionally, these experiments include flexible deployment, continuing the investigation from Experiment 2, but this time exploring how policy and flexible decision-making work together. For instance, the option to add a warehouse was never exercised in the Experiment 2 cases, but with an

external catalyst such as policy incentives, this option may be triggered. This experiment investigates whether such policy-triggered investments can reduce the economic barriers to achieving greater sustainability and circularity.

This addresses gaps in literature because it tests the partnership between policy and CONOPS never before proposed, creating a computational laboratory to assess policy mechanisms and their varying parameters across various versions of the future to analyze policy for robustness. Previously, no research has considered or attempted to quantify the benefit that different policies could have on the business case for LEO-based OOS and how they may affect sustainability metrics like reduced emissions and increased refurbishment. Existing policy analyses for space sustainability have remained largely qualitative or focused on compliance costs, rather than examining how policy instruments can enable economically viable infrastructure that simultaneously improves environmental outcomes.

These experiments sought to identify what combination of policy, flexibility, and initial investments provide the best sustainability benefits for the least expense to both satellite operator and government, recognizing that there may be tradeoffs depending on stakeholder priorities. The framework provides a tool for better understanding those tradeoffs and their conditions.

4.4.1 Experiment 3.1: Policy Parameter Down-Selection Analysis

This section presents a comprehensive sensitivity analysis to identify high-performing parameter configurations for eight policy schemes. The experiment employed variance-based global sensitivity analysis across 20 scenarios per configuration, which was a deliberate choice to achieve sufficient ranking stability for identifying top contenders while managing computational expense. This scenario count provided reliable first-place rankings for down-selection purposes, with the understanding that selected configurations would undergo more extensive evaluation in subsequent experiments.

Performance was evaluated through multiple objectives: total customer costs, total emissions, number of refurbishments, and a weighted composite metric balancing economic and environmental trade-offs. To manage the design space, satellite upgrade configuration (Rf) and warehouse upgrade (Upg) variables were set to 1 (flexible).

At the end of each scenario, the government subsidy balances were subtracted from total customer costs to properly account for net expenditure, ensuring fair comparison across policies. This treats the initial subsidy allocation as a government loan reconciled at the 30-year endpoint. Policy 4 is excluded from this treatment as premiums are collected by a private firm.

Statistical significance testing distinguished between `policy_scheme_config` significance (indicating policy effectiveness versus no policy) and individual parameter significance (indicating parameter value importance within the policy mechanism). This analysis yielded 18 down-selected configurations across all eight policies for subsequent comparative evaluation. Table 4.18 provides a summary of Experiment 3.1 results.

Table 4.17: Policy Configuration Parameters

Abbrev.	Parameter	Abbrev.	Parameter
REBATE	rebate_for_refurbishment	TAX_PCT	tax_percentage
INIT_SUB	policy_scheme_initial_sub	TAX_SHAPE	tax_shape_parameter
OUF	annualOUF	FINE	fine_config
REFUND	refund_condition	POLICY_EN	policy_scheme_config
PREM	annual_premium		

Configuration Selection Rationale

Each policy yielded 1-3 configurations with distinct performance profiles. Configuration x.1 variants consistently prioritized cost-neutrality with moderate parameters. Configuration x.2 variants emphasized environmental performance through more aggressive settings, accepting higher costs for superior emissions outcomes. Configuration x.3 variants (where present) explored intermediate positions or specialized

objectives such as increased refurbishment throughput.

Policy 1 (OUF with Refund): OUF was significant for all cost and emissions percentiles. Configuration 1.1 (OUF=\$50k, no rebate/subsidy, refund=1) achieved cost-neutrality. Configuration 1.2 (OUF=\$75k) maximized emissions reduction while accepting higher costs.

Policy 2 (Subsidy with Rebate): POLICY_EN was significant for average emissions, while individual parameters affected only extreme percentiles. Configuration 2.1 (REBATE=\$250k, INIT_SUB=1, OUF=\$2.5k) established a cost-neutral configuration. Configuration 2.2 (OUF=\$25k) pursued deorbit reductions through tenfold OUF increase.

Policy 3 (Fine-Based): POLICY_EN was significant for average emissions; FINE affected 95th percentile only. Configuration 3.1 (REBATE=\$250k, INIT_SUB=1, FINE=\$100k) balanced cost and emissions. Configuration 3.2 (FINE=\$10M) maximized enforcement through severe penalties for worst-case scenarios.

Policy 4 (Annual Premium): POLICY_EN was significant for average and 5th percentile emissions, but premium rates showed limited parameter significance. Configuration 4.1 (PREM=0.001) achieved superior cost-weighted performance, reflecting that benefits derive from the policy mechanism itself rather than premium intensity.

Policy 5 (Tax-Based): REBATE was significant across all cost percentiles, making it critical for this policy. TAX_PCT affected average and 5th percentile emissions. Configuration 5.1 (REBATE=\$250k, TAX_PCT=0.001) achieved best cost-weighted balance. Configuration 5.2 (TAX_PCT=0.005) prioritized deorbit reductions through aggressive taxation. Configuration 5.3 (no rebate, TAX_PCT=0.003) explored intermediate emissions-focused approach.

Policy 6 (Tax and Fine): TAX_PCT was significant for average and 95th percentile emissions; FINE affected 5th and 95th percentiles. Configuration 6.1 (REBATE=\$250k, INIT_SUB=1, TAX_PCT=0.001, FINE=\$500k) leveraged complementary enforcement: taxes for average performance, fines for extreme cases.

Policy 7 (OUF-Fine): This policy had the broadest observed significance patterns. All parameters were significant for total cost (except REBATE at 95th percentile). OUF, REBATE, and FINE were significant for emissions at the 5th and 95th percentiles. Configuration 7.1 (REBATE=\$250k, INIT_SUB=1, FINE=\$500k, OUF=\$2.5k) established a cost-focused configuration with synergistic OUF-fine combination. Configuration 7.2 (FINE=\$1M, OUF=\$5k) targeted extreme percentile emissions. Configuration 7.3 (FINE=\$10M, OUF=\$25k) focused on reducing deorbit through aggressive dual enforcement.

Policy 8 (Premium with Superfund): Unlike Policy 4, this policy's collected premiums contribute to a subsidy fund. REBATE was significant for average costs; PREM affected only 5th percentile emissions. Configuration 8.1 (REBATE=\$250k, INIT_SUB=1, PREM=0.001) represented a cost-superior configuration. Configuration 8.2 (PREM=0.02) maximized refurbishment throughput through elevated premium rates that feed a larger subsidy fund pool.

Key Findings and Implications

Diminishing Returns from Aggressive Interventions: Analysis consistently revealed that moderate parameter values outperformed extreme settings for cost performance, as expected. Pushing fees or fines beyond certain thresholds yielded diminishing marginal benefits, adding costs without proportional sustainability performance improvements.

Complementary Enforcement Mechanisms: Combined policies (particularly Policies 6 and 7) demonstrated that pairing continuous economic incentives (taxes, OUF) with episodic penalties (fines) effectively manage the full performance distribution. Taxes shape average behavior while fines address extreme cases, creating robust enforcement coverage.

Parameter Significance Patterns: Refurbishment rebates showed significance at 95th percentiles across multiple policies, with Policy 5 demonstrating particularly broad rebate effectiveness. Initial subsidies primarily moderated emissions variability at extremes rather than shifting central tendencies. Enforcement parameters exhibited distinct roles: fines managed extremes (5th/95th percentiles), taxes affected average and extreme performance depending on pairing, OUF demonstrated robust significance across distributions, and premium rates showed limited parameter significance where policy mechanism itself drove benefits.

Policy Mechanism vs. Parameter Intensity: Several policies (notably 3, 4, and 8) demonstrated POLICY_EN significance where specific parameters did not, revealing that enabling these policies matters more than fine-tuning parameter values within tested ranges. This contrasts with Policies 2, 5, 6, and 7, where careful parameter calibration significantly affects outcomes.

Trade-Off Mapping: The tiered configuration approach (x.1, x.2, x.3 variants) helped map the cost-priority versus performance trade-off space. Configuration selection provides policymakers with clear choices between maintaining cost-neutrality with modest gains or accepting higher costs for aggressive environmental outcomes, with quantified trade-offs for each option.

Table 4.18: Experiment 3.1 Down-Selection Summary

Policy Scheme	Description	Parameters Tested	Config	Down-Selected Settings	Focus/Objective
Policy 1	OUF with Refund Conditions	OUF={10k, 25k, 50k, 75k}, REBATE={0, 250k}, INIT_SUB={0, 1}, REFUND={1, 2}	1.1	OUF=50k, REBATE=0, INIT_SUB=0, REFUND=1	Cost-neutral baseline
			1.2	OUF=75k, REBATE=0, INIT_SUB=0, REFUND=1	Environmental (max emissions reduction)
Policy 2	OUF, Initial Subsidy with Rebate	REBATE={0, 250k}, INIT_SUB={0, 1}, OUF={2.5k, 5k, 10k, 25k}	2.1	OUF=2.5k, REBATE=250k, INIT_SUB=1	Cost-neutral baseline
			2.2	OUF=25k, REBATE=250k, INIT_SUB=1	Reduced deorbits
Policy 3	Fine-Based Enforcement	REBATE={0, 250k}, INIT_SUB={0, 1}, FINE={100k, 500k, 1M, 10M}	3.1	FINE=100k, REBATE=250k, INIT_SUB=1	Cost-neutral with competitive emissions
			3.2	FINE=10M, REBATE=250k, INIT_SUB=1	Focus on deorbit reduction
Policy 4	Annual Premium	REBATE={0}, INIT_SUB={0}, PREM={0.001, 0.005, 0.01, 0.025, 0.05, 0.1}	4.1	PREM=0.001, REBATE=0, INIT_SUB=0	Minimal intervention approach
Policy 5	Tax-Based Mechanism	REBATE={0, 250k}, INIT_SUB={0, 1}, TAX_PCT={0.001, 0.003, 0.005, 0.01}	5.1	TAX_PCT=0.001, REBATE=250k, INIT_SUB=0	Cost-neutral with efficiency
			5.2	TAX_PCT=0.005, REBATE=250k, INIT_SUB=0	Reduced deorbits
			5.3	TAX_PCT=0.003, REBATE=0, INIT_SUB=0	Intermediate emissions focus
Policy 6	Combined Tax & Fine	REBATE={0, 250k}, INIT_SUB={0, 1}, TAX_PCT={0.001, 0.003, 0.005, 0.01}, FINE={100k, 500k, 1M, 10M}	6.1	TAX_PCT=0.001, FINE=500k, REBATE=250k, INIT_SUB=1	Cost-neutral with balanced enforcement
Policy 7	OUF-Fine Combination	REBATE={0, 250k}, INIT_SUB={0, 1}, FINE={100k, 500k, 1M, 10M}, OUF={2.5k, 5k, 10k, 25k}	7.1	OUF=2.5k, FINE=500k, REBATE=250k, INIT_SUB=1	Cost-priority configuration
			7.2	OUF=5k, FINE=1M, REBATE=250k, INIT_SUB=1	Emissions performance at extremes
			7.3	OUF=25k, FINE=10M, REBATE=250k, INIT_SUB=1	Reduced deorbits
Policy 8	Premium-Based Incentive	REBATE={0, 250k}, INIT_SUB={0, 1}, PREM={0.001, 0.005, 0.01, 0.025, 0.05, 0.1}	8.1	PREM=0.001, REBATE=250k, INIT_SUB=1	Cost-neutral with subsidy fund
			8.2	PREM=0.02, REBATE=250k, INIT_SUB=1	Maximize refurbishment activity

These 18 down-selected configurations advance to Experiment 3.2 for comprehensive head-to-head comparison under identical simulation conditions with larger scenario batches. The down-selection approach enabled efficient screening of the vast parameter space while ensuring that top-performing configurations from each policy scheme receive rigorous large-scale evaluation. Experiment 3.2 incorporates infrastructure upgrade pathway variations (Upg) to assess whether flexible versus inflexible warehouse upgrade strategies affect relative policy performance across the tested variants.

4.4.2 Experiment 3.2: Comparison and Analysis for Down-Selected Policies

This section presents comprehensive results of the Value at Risk and Gain (VARG) analysis combined with statistical significance testing for the selection of policies from Experiment 3.1 that show promise in improving infrastructure feasibility or efficiency through either cost-neutral policies or more ambitious policies that seek to alter infrastructure composition. Table 4.18 describes the down-selected policy combinations and Table 4.19 details the infrastructure configuration variations.

Configuration Set

Table 4.19: Infrastructure Configurations for Policy Comparison Experiment 3.2 and 3.3

Parameter	Values	Description
depot_config	0, 1	0: No depot/warehouse infrastructure; 1: Depot enabled (prerequisite for all other advanced configurations)
CAAS_config	1	1: CAAS enabled (only when depot_config = 1)

Parameter	Values	Description
refuelable_config	1	1: Refuelable satellites option available (only when depot_config = 1)
upgraded_config	0, 1, 2	0: No warehouse upgrades; 1: Upgraded warehouse option available (only when depot_config = 1 and refuelable_config = 1); 2: Warehouse upgrades immediately available
add_new_wh_config	1	1: Allow new warehouse additions (only when depot_config = 1)
rpo_config	2	2: Satellite RPO immediately available (only when depot_config = 1)

Total configurations per scenario:

- Baseline scenario: 1 configuration (`depot_config` = 0)
- Depot scenarios: 17 policy variants \times 3 upgrade variants = 51 configurations
- **Total: 52 configurations per scenario replication**

The experiment design ensures that each policy is evaluated under identical infrastructure conditions (when depot is enabled), with the upgrade flexibility parameter (`upgraded_config`) providing insight into how different satellite upgrade strategies interact with policy mechanisms. This configuration enables direct, fair comparison of policy effectiveness while controlling for infrastructure and operational uncertainties.

VARG Results for Experiment 3.2

The results for Experiment 3.2 are provided in Table 4.21 through Table 4.23. For experiment 3.2, the top-ranked CAAS+policy configuration after 80 scenarios is a policy 3 configuration with 1 depot (D:1, CAAS:1, Rf:1, Upg:0, P:3, RPO:2, Fine:100km,

Table 4.20: Top Performing Configuration Parameter Settings - Experiment 3.2

Configuration	Depot	Rf	Upg	RPO	POLICY_EN	REBATE	FINE	INIT_SUB	OUF
0-Warehouse Baseline	0	0	0	0	0	0	0	0	0
Best Cost No-Policy	1	1	0	2	0	0	0	0	0
Best for Total Cost	1	1	0	2	3	250k	100k	0	0
Best for Total Emissions & Weighted Multi-Objective	1	1	2	2	1	0	0	0	75k
Best for Cost/Refurbish, Cost/Collection	1	1	0	2	7	250k	10M	1	25k
Best for Sats Refurb (Space)	1	1	2	2	7	250k	10M	1	25k
Best for Num Sat Refurb	1	1	1	2	7	250k	10M	1	25k

Rebate: 250k). However, the 0-warehouse baseline, ranked 1st, is the only stable ranking for the given number of scenarios. Furthermore, this configuration's 0.43% increase in total cost compared to the 0-warehouse baseline is too small to be statistically significant. The only policy configurations present in the top performing configurations to have a significant cost increase over the 0-warehouse baseline and top no-policy CAAS configuration are the policy 7 configurations; otherwise, cost is comparable. Additionally, none of the identified policies offer statistically significant improvements in emissions compared to the top-cost no-policy configuration.

The results for Experiment 3.2 are provided in Table 4.21 through Table 4.23. For Experiment 3.2, the top-ranked CAAS+policy configuration after 80 scenarios is a Policy 3 configuration with 1 depot (D:1, CAAS:1, Ref:1, Upg:0, P:3, RPO:2, Fine:100k, Rebate:250k). However, the 0-warehouse baseline, ranked 1st, is the only stable ranking for the given number of scenarios. Furthermore, this configuration's 0.43% increase in total cost compared to the 0-warehouse baseline, depicted in Figure 4.10, is too small to be statistically significant. The only policy configurations present in the top performing configurations to have a significant cost increase over the 0-warehouse baseline and top no-policy CAAS configuration are the Policy 7 configurations; otherwise, cost is comparable. Additionally, none of the identified policies offer statistically significant improvements in emissions compared to the top-cost no-policy configuration.

The top emissions policy configuration, depicted in Figure 4.11, demonstrates com-

parable cost to both the 0-warehouse baseline and the top no-policy configuration. It achieves a 7.45% reduction in emissions compared to the 0-warehouse baseline, which is statistically significant compared to the 0-warehouse baseline ($p < 0$, 95% CI: [3.14M, 3.31M], MoE: 2.66%). While this reduction represents an improvement over the top no-policy configuration, the difference between these two configurations is not statistically measurable given the number of cases ($p = 0.56$). However, this same emissions configuration demonstrates measurably better numbers of refurbished satellites compared to both the 0-warehouse baseline and the top no-policy configuration, with a 29.47% increase over the top-cost no-policy configuration ($p = 0.038$, 95% CI: [65, 93], MoE: 16.9%). This is a key distinction from the findings in Experiment 2, where the emissions-superior configuration (with the same refuel and upgrade settings) did not offer statistically more refurbishments than the non-flexible option.

These results demonstrate how the model was able to cut through the noise and identify a policy that provides sustainability benefits. While the policy's emissions improvements were too small to measure compared to the no-policy configuration given the number of cases, it was able to provide greater, measurable rates of refurbishment compared to the no-policy configuration identified in Experiment 2. This shows how the model could find a cost-neutral policy capable of improving sustainability metrics without measurable increases in the total cost over the 30-year simulation period.

Experimental Observation 3.1

The \$75,000 OUF scheme that provides a rebate when satellites are collected provides a cost-neutral policy approach that improves refurbishment throughput relative to the top no-policy configuration.

The cost-superior CAAS+policy configuration never gains additional warehouses. If the priority is to grow the OOS infrastructure over the 30-year horizon, variations of policies 2 and 3 provide the means of expansion, albeit with penalties to total cost.




Figure 4.10: Experiment 3.2: Total Cost VARG Plot




Figure 4.11: Experiment 3.2: Total NOx Emissions VARG Plot

Over the simulated 30-year period, policy (D:1, CAAS:1, Ref:1, Upg:0, P:2, RPO:2, OUF:2.5e4, Psub:1, Rebate: 2.5e5) creates an average of 7.5 warehouses and has an avg total cost that is 11.47% higher than the 0-warehouse baseline. In comparison, the no-policy configuration that starts with 7 warehouses (D:7, CAAS:1, Ref:1, Upg:0, P:0, RPO:2) is 16.221% more expensive than the 0-warehouse baseline (Welch's t-test p-value<0.001). Similarly, policy (D:1, CAAS:1, Ref:1, Upg:0, P:3, RPO:2, Fine:1e7, Psub:1, Rebate: 2.5e5) creates 4 warehouses on average with an average total cost that is 5.292% greater than the 0-warehouse baseline, compared to the no-policy configuration that starts with 4 warehouses (D:4, CAAS:1, Ref:1, Upg:0, P:0, RPO:2) that has a 7.59% greater average total cost (Welch's t-test p-value=0.0286). Using these growth-oriented policies creates a CAAS infrastructure of comparable size over the 30 year timeline, leveraging the time value of money to reduce total cost. However, this approach comes with a trade-off, since the grow-to-7 and grow-to-4 configurations have less emissions benefits and fewer refurbished satellites compared to the initialized 7-warehouse and 4-warehouse configurations. Table 4.21 through Table 4.23 provide a breakdown of how the growth policies compare to the large, no-policy configurations in terms of total cost, emissions, and refurbishments.

Experimental Observation 3.2

Policy+flexible warehouse roll-out enables the framework to reach a certain warehouse number for less total cost compared to initializing the infrastructure with more warehouses, but comes with less capability to improve emissions and refurbishments over the 30 years. To make real change in emissions over the 30-year horizon, there need to be more warehouses at the start of the simulation.

Table 4.21: Experiment 3.2: Total Cost

Top Performing Metric	Configuration	Percentile	Total Cost Value (\$)	Total Cost CI Lower	Total Cost CI Upper	Total Cost MoE%	Total Cost % vs Baseline	Total Cost p-value (Baseline)	Total Cost Sig? (Baseline)	Total Cost % vs Best No-Policy	Total Cost p-value (No-Policy)	Total Cost Sig? (No-Policy)
0-Warehouse Baseline	d:0 up:0 rf:0 rpo:0 P:0 OUF:0 Fine:0 Rebate:0	avg	4.46E+09	4.38E+09	4.55E+09	1.905	N/A	N/A	N/A	-0.402	0.758	No
		5th	3.98E+09	3.93E+09	4.09E+09	2.025	N/A	N/A	N/A	-2.074	0.758	No
		95th	5.12E+09	5.01E+09	5.34E+09	3.232	N/A	N/A	N/A	0.002	0.758	No
Best Cost No-Policy	d:1 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	4.48E+09	4.40E+09	4.56E+09	1.773	0.404	0.758	No	N/A	N/A	N/A
		5th	4.06E+09	4.03E+09	4.10E+09	0.808	2.118	0.758	No	N/A	N/A	N/A
		95th	5.12E+09	4.97E+09	5.33E+09	3.557	-0.002	0.758	No	N/A	N/A	N/A
Best for Total Cost	d:1 up:0 rf:1 rpo:2 P:3 OUF:2.5e4 Fine:100k Rebate:250k	avg	4.48E+09	4.40E+09	4.56E+09	1.799	0.437	0.741	No	0.033	0.979	No
		5th	4.05E+09	4.04E+09	4.09E+09	0.654	1.788	0.741	No	-0.323	0.979	No
		95th	5.12E+09	4.91E+09	5.35E+09	4.298	0.008	0.741	No	0.010	0.979	No
Best for Total Emissions & Weighted Multi-Objective	d:1 up:2 rf:1 rpo:2 P:1 OUF:7.5e4 R:1 Fine:0 Rebate:0	avg	4.55E+09	4.45E+09	4.65E+09	2.246	1.926	0.200	No	1.516	0.298	No
		5th	4.08E+09	4.08E+09	4.13E+09	0.633	2.626	0.200	No	0.497	0.298	No
		95th	5.20E+09	4.98E+09	5.93E+09	9.151	1.557	0.200	No	1.559	0.298	No
Best for Cost/Refurb, Cost/Collect	d:1 up:0 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	5.26E+09	5.18E+09	5.35E+09	1.691	17.998	0.000	Yes	17.523	0.000	Yes
		5th	4.77E+09	4.69E+09	4.81E+09	1.255	19.974	0.000	Yes	17.486	0.000	Yes
		95th	5.90E+09	5.77E+09	6.04E+09	2.267	15.120	0.000	Yes	15.122	0.000	Yes
Best for Sats Refurb. In Space	d:1 up:2 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	5.37E+09	5.27E+09	5.46E+09	1.766	20.297	0.000	Yes	19.813	0.000	Yes
		5th	4.84E+09	4.79E+09	4.88E+09	0.934	21.717	0.000	Yes	19.193	0.000	Yes
		95th	6.09E+09	5.86E+09	6.15E+09	2.386	18.904	0.000	Yes	18.906	0.000	Yes
Best for Total Num Sat Refurb	d:1 up:1 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	5.32E+09	5.23E+09	5.41E+09	1.681	19.228	0.000	Yes	18.749	0.000	Yes
		5th	4.80E+09	4.75E+09	4.89E+09	1.424	20.716	0.000	Yes	18.212	0.000	Yes
		95th	5.96E+09	5.79E+09	6.11E+09	2.719	16.370	0.000	Yes	16.373	0.000	Yes
Grow-to-7 Warehouses	d:1 up:0 rf:1 rpo:2 P:2 OUF:2.5e4 Fine:0 Rebate:250k	avg	4.97E+09	4.89E+09	5.05E+09	1.613	11.469	0.000	Yes	11.020	0.000	Yes
		5th	4.52E+09	4.40E+09	4.58E+09	1.984	13.765	0.000	Yes	11.405	0.000	Yes
		95th	5.61E+09	5.41E+09	5.75E+09	3.010	9.547	0.000	Yes	9.549	0.000	Yes
Grow-to-4 Warehouses	d:1 up:0 rf:1 rpo:2 P:3 OUF:0 Fine:10M Rebate:250k	avg	4.70E+09	4.61E+09	4.79E+09	1.921	5.292	0.000	Yes	4.868	0.000	Yes
		5th	4.25E+09	4.22E+09	4.26E+09	0.500	6.892	0.000	Yes	4.676	0.000	Yes
		95th	5.40E+09	5.21E+09	5.53E+09	2.952	5.382	0.000	Yes	5.384	0.000	Yes
7 Initial Warehouses, no policy	d:7 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	5.19E+09	5.06E+09	5.31E+09	2.489	16.221	0.000	Yes	15.754	0.000	Yes
		5th	4.77E+09	4.70E+09	4.85E+09	1.529	20.081	0.000	Yes	17.590	0.000	Yes
		95th	5.84E+09	5.50E+09	6.19E+09	5.938	13.948	0.000	Yes	13.951	0.000	Yes
4 Initial Warehouses, no policy	d:4 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	4.80E+09	4.67E+09	4.93E+09	2.766	7.593	0.000	Yes	7.160	0.000	Yes
		5th	4.42E+09	4.38E+09	4.44E+09	0.640	11.090	0.000	Yes	8.786	0.000	Yes
		95th	5.49E+09	5.15E+09	5.89E+09	6.777	7.192	0.000	Yes	7.194	0.000	Yes

Table 4.22: Experiment 3.2: Total NOx Emissions

Top Performing Metric	Configuration	Percentile	Total Emissions Value	Total Emissions CI Lower	Total Emissions CI Upper	Total Emissions MoE%	Total Emissions % vs Baseline	Total Emissions p-value (Baseline)	Total Emissions Sig? (Baseline)	Total Emissions % vs Best No-Policy	Total Emissions p-value (No-Policy)	Total Emissions Sig? (No-Policy)
0-Warehouse Baseline	d:0 up:0 rf:0 rpo:0 P:0 OUF:0 Fine:0 Rebate:0	avg	3.48E+06	3.41E+06	3.56E+06	2.083	N/A	N/A	N/A	6.860	0.000	Yes
		5th	2.90E+06	2.83E+06	3.04E+06	3.585	N/A	N/A	N/A	14.048	0.000	Yes
		95th	3.94E+06	3.82E+06	3.98E+06	1.936	N/A	N/A	N/A	3.574	0.000	Yes
Best Cost No-Policy	d:1 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	3.26E+06	3.17E+06	3.35E+06	2.748	-6.419	0.000	Yes	N/A	N/A	N/A
		5th	2.54E+06	2.32E+06	2.75E+06	8.564	-12.318	0.000	Yes	N/A	N/A	N/A
		95th	3.81E+06	3.71E+06	3.87E+06	2.105	-3.451	0.000	Yes	N/A	N/A	N/A
Best for Total Cost	d:1 up:0 rf:1 rpo:2 P:3 OUF:2.5e4 Fine:100k Rebate:250k	avg	3.26E+06	3.17E+06	3.35E+06	2.793	-6.416	0.000	Yes	0.003	0.999	No
		5th	2.50E+06	2.36E+06	2.71E+06	7.026	-13.554	0.000	Yes	-1.410	0.999	No
		95th	3.83E+06	3.69E+06	3.97E+06	3.611	-2.989	0.000	Yes	0.478	0.999	No
Best for Total Emissions & Weighted Multi-Objective	d:1 up:2 rf:1 rpo:2 P:1 OUF:7.5e4 R:1 Fine:0 Rebate:0	avg	3.22E+06	3.14E+06	3.31E+06	2.662	-7.446	0.000	Yes	-1.097	0.567	No
		5th	2.52E+06	2.38E+06	2.64E+06	5.061	-13.070	0.000	Yes	-0.858	0.567	No
		95th	3.74E+06	3.67E+06	3.78E+06	1.481	-5.079	0.000	Yes	-1.687	0.567	No
Best for Cost/Refurb, Cost/Collect	d:1 up:0 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	3.30E+06	3.24E+06	3.37E+06	2.072	-5.114	0.001	Yes	1.396	0.423	No
		5th	2.70E+06	2.54E+06	2.89E+06	6.603	-6.625	0.001	Yes	6.492	0.423	No
		95th	3.73E+06	3.66E+06	3.77E+06	1.588	-5.330	0.001	Yes	-1.947	0.423	No
Best for Sats Refurb. In Space	d:1 up:2 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	3.33E+06	3.26E+06	3.40E+06	2.111	-4.402	0.003	Yes	2.156	0.221	No
		5th	2.72E+06	2.65E+06	2.86E+06	3.860	-6.088	0.003	Yes	7.105	0.221	No
		95th	3.76E+06	3.68E+06	3.92E+06	3.243	-4.575	0.003	Yes	-1.165	0.221	No
Best for Total Num Sat Refurb	d:1 up:1 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	3.27E+06	3.20E+06	3.35E+06	2.226	-6.022	0.000	Yes	0.424	0.812	No
		5th	2.65E+06	2.57E+06	2.89E+06	6.068	-8.459	0.000	Yes	4.401	0.812	No
		95th	3.72E+06	3.64E+06	3.80E+06	2.167	-5.710	0.000	Yes	-2.340	0.812	No
Grow-to-7 Warehouses	d:1 up:0 rf:1 rpo:2 P:2 OUF:2.5e4 Fine:0 Rebate:250k	avg	3.25E+06	3.17E+06	3.33E+06	2.382	-6.673	0.000	Yes	-0.271	0.882	No
		5th	2.67E+06	2.56E+06	2.78E+06	4.196	-7.629	0.000	Yes	5.347	0.882	No
		95th	3.82E+06	3.64E+06	3.88E+06	3.180	-3.131	0.000	Yes	0.332	0.882	No
Grow-to-4 Warehouses	d:1 up:0 rf:1 rpo:2 P:3 OUF:0 Fine:10M Rebate:250k	avg	3.27E+06	3.18E+06	3.35E+06	2.516	-6.207	0.000	Yes	0.227	0.904	No
		5th	2.62E+06	2.44E+06	2.73E+06	5.540	-9.639	0.000	Yes	3.055	0.904	No
		95th	3.74E+06	3.70E+06	3.80E+06	1.330	-5.171	0.000	Yes	-1.782	0.904	No
7 Initial Warehouses, no policy	d:7 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	3.08E+06	2.98E+06	3.18E+06	3.301	-11.648	0.000	Yes	-5.587	0.008	Yes
		5th	2.54E+06	2.39E+06	2.80E+06	8.079	-12.360	0.000	Yes	-0.048	0.008	Yes
		95th	3.48E+06	3.32E+06	3.53E+06	2.951	-11.637	0.000	Yes	-8.479	0.008	Yes
4 Initial Warehouses, no policy	d:4 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	3.13E+06	3.02E+06	3.25E+06	3.634	-10.052	0.000	Yes	-3.882	0.082	No
		5th	2.56E+06	2.42E+06	2.80E+06	7.275	-11.560	0.000	Yes	0.864	0.082	No
		95th	3.60E+06	3.40E+06	3.70E+06	4.168	-8.619	0.000	Yes	-5.352	0.082	No

Table 4.23: Experiment 3.2: Number of Refurbishments

Top Performing Metric	Configuration	Percentile	Num Sat Refurb Value	Num Sat Refurb CI Lower	Num Sat Refurb CI Upper	Num Sat Refurb MoE%	Num Sat Refurb % vs No-Policy	Num Sat Refurb p-value (No-Policy)	Num Sat Refurb Sig? (No-Policy)
0-Warehouse Baseline	d:0 up:0 rf:0 rpo:0 P:0 OUF:0 Fine:0 Rebate:0	avg	0.00	0.00	0.00	-	N/A	N/A	N/A
		5th	0.00	0.00	0.00	-	N/A	N/A	N/A
		95th	0.00	0.00	0.00	-	N/A	N/A	N/A
Best Cost No-Policy	d:1 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	61.16	50.37	71.95	17.643	N/A	N/A	N/A
		5th	0.00	0.00	0.00	-	N/A	N/A	N/A
		95th	145.70	125.00	164.00	13.384	N/A	N/A	N/A
Best for Total Cost	d:1 up:0 rf:1 rpo:2 P:3 OUF:2.5e4 Fine:100k Rebate:250k	avg	61.66	51.19	72.14	16.985	0.817	0.947	No
		5th	0.00	0.00	3.90	-	-	-	No
		95th	159.35	110.00	176.00	20.709	9.369	0.947	No
Best for Total Emissions & Weighted Multi-Objective	d:1 up:2 rf:1 rpo:2 P:1 OUF:7.5e4 R:1 Fine:0 Rebate:0	avg	79.19	65.80	92.58	16.907	29.471	0.039	Yes
		5th	29.70	9.00	38.00	48.822	inf	0.039	Yes
		95th	140.20	128.15	175.00	16.708	-3.775	0.039	Yes
Best for Cost/Refurb, Cost/Collect	d:1 up:0 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	273.10	252.82	293.38	7.426	346.515	0.000	Yes
		5th	114.50	63.00	170.00	46.725	inf	0.000	Yes
		95th	403.20	387.55	444.00	7.000	176.733	0.000	Yes
Best for Sats Refurb, In Space	d:1 up:2 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	244.09	221.47	266.70	9.265	299.080	0.000	Yes
		5th	75.70	44.00	118.85	49.439	inf	0.000	Yes
		95th	401.30	369.30	433.00	7.937	175.429	0.000	Yes
Best for Total Num Sat Refurb	d:1 up:1 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	283.34	262.34	304.34	7.412	363.254	0.000	Yes
		5th	66.95	57.00	156.15	74.050	inf	0.000	Yes
		95th	401.95	381.15	445.00	7.943	175.875	0.000	Yes
Grow-to-7 Warehouses	d:1 up:0 rf:1 rpo:2 P:2 OUF:2.5e4 Fine:0 Rebate:250k	avg	193.13	178.42	207.83	7.616	215.757	0.000	Yes
		5th	88.15	69.00	109.75	23.114	inf	0.000	Yes
		95th	298.30	273.10	333.00	10.040	104.736	0.000	Yes
Grow-to-4 Warehouses	d:1 up:0 rf:1 rpo:2 P:3 OUF:0 Fine:10M Rebate:250k	avg	119.81	104.57	135.06	12.725	95.892	0.000	Yes
		5th	34.70	24.00	46.80	32.853	inf	0.000	Yes
		95th	228.10	209.00	256.00	10.302	56.555	0.000	Yes
7 Initial Warehouses, no policy	d:7 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	235.65	187.13	284.17	20.590	285.280	0.000	Yes
		5th	0.65	0.00	84.95	6534.615	inf	0.000	Yes
		95th	425.05	381.35	492.00	13.016	191.730	0.000	Yes
4 Initial Warehouses, no policy	d:4 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	198.18	161.08	235.27	18.719	224.016	0.000	Yes
		5th	18.85	0.00	89.81	238.223	inf	0.000	Yes
		95th	344.75	282.45	422.00	20.239	136.616	0.000	Yes

Overall, comparing these two expansion-promoting policies with no-policy configurations reveals the cost benefits of flexible warehouse deployment. This observation motivates the following experimentation set, which varies the combinations of down-selected policies with the number of initial warehouses.

4.4.3 Experiment 3.3: Policy and Flexibility for Configurations with Greater Than 1 Initial Depot

Experiment 3.3 takes a further down-selection of policies from experiment 3.2 (policies 1 and 7) and expand the design space to enable policy for 1, 4, or 7 initial warehouses as well as compare flexibly upgraded satellites with pre-initialized satellites upgrades (Rf 1 to Rf 1,2). Since the number of depots is a driving contributor of emissions reduction, this experiment reveals the potential for policy to enable a larger number of initial depots in a manner that is more financially palatable than would be possible without government intervention. Table 4.24 contains the details on the down-selected policy configurations. Over 70 uncertain scenarios, the infrastructure is varied in the same manner as before, as detailed in Table 4.19.

Configuration Set

Table 4.24: Downselected Parameter Configurations for Policy Comparison Experiment 3.3

Policy Scheme	Rebate Refurb	Initial Sub	OUF	Rebate Cond.	Premium	Tax %	Tax Shape	Fine
0 (Baseline)	0	0	0	0	0	0	0	0
1.1	0	0	50k	0	0	0	0	0
1.2	0	0	75k	0	0	0	0	0
7.1	250k	1	2.5k	0	0	0	0	500k
7.2	250k	1	5k	0	0	0	0	1M
7.3	250k	1	25k	0	0	0	0	10M

Bold values indicate active policy parameters. k = thousand, M = million.

VARG Results for Experiment 3.3

Comparing the various configurations in experiment 3.3, the configuration with 4 initial depots and policy scheme 1 (D:4, CAAS:1, Ref:1, Upg:0, P:1, RPO:2, OUF:5e4) provides the best weighted multi-objective result cost compared to other policy configurations when the priority is 60/40 for total cost vs total emissions. Relative to the 0-warehouse baseline, this configuration shows a +7.80% [95% CI: [4.78B,4.93B], %MoE:1.52] increase in total average cost and -9.94% decrease in average emissions [95% CI: [3.06M,3.21M], %MoE:2.36]. The difference in cost is statistically significant compared to both the 0-warehouse baseline ($p<0.001$) and top cost no-policy configuration ($p<0.001$ for D:1, CAAS:1, Ref:1, Upg:0, P:0, RPO:2) while the difference in emissions is statistically significant for 0-warehouse baseline ($p<0.001$) but marginally significant for the top cost no-policy configuration ($p=0.073$).

This configuration demonstrates sustainability improvements over its 1-depot variant, (D:1, CAAS:1, Ref:1, Upg:0, P:1, RPO:2, OUF:5e4), achieving greater average emission reductions (-9.94% decrease vs. -6.86% with $p=0.0058$ from Welch's t-test) while refurbishing substantially more satellites (262 vs. 72 with $p<0.001$). These sustainability benefits do come at a cost however, with the best weighted multi-objective configurations costing 7.81% more than the 0-warehouse baseline and the 1-depot variant costing about the same as the 0-warehouse baseline.

Table 4.25: Top Performing Configuration Parameter Settings - Experiment 3.3

Configuration	Depot	Active	Rf	Upg	Pol	RPO	Rebate	Fine	Sub	OUF	CAAS
0-Warehouse Baseline	0	0	0	0	0	0	0	0	0	0	0
Best Cost No-Policy	1	1	1	0	0	2	0	0	0	0	0
Best for Total Cost	1	1	1	0	1	2	0	0	0	75k	1
Best for Emissions & Best Emissions Savings/Cost	7	1	1	0	1	2	0	0	0	50k	1
Best for Cost/Refurbish	4	1	1	1	7	2	250k	10M	1	25k	0
Best for Sats Refurb. In Space	7	1	1	2	7	2	250k	10M	1	25k	0
Best for Total Num Sat Refurb & Cost/Collection	7	1	1	0	7	2	250k	10M	1	25k	0
Best for Weighted Multi-Objective	4	1	1	0	1	2	0	0	0	50k	1

This policy configuration also outperforms its no-policy equivalent (D:4, CAAS:1, Ref:1, Upg:0, P:0) in terms of sustainability. While the policy version shows only slight differences in emissions and cost (7.80% vs. 7.82% increase in average cost compared to the 0-warehouse baseline; -9.94% vs. -8.81% decrease in emissions compared to the 0-warehouse baseline with comparative p-value=0.3 from Welch's t-test), it delivers far more average total refurbishments (262 versus 194 with p<0.001) and achieves a lower cost per refurbishment (-95.79% vs. -79.45% decrease compared to top cost no-policy configuration). This demonstrates how the same cost-neutral policy that encouraged greater throughput in the 1-depot variation provides similar benefits for the 4-depot configuration. Starting with 4 depots yields greater sustainability benefits while the policy enhances throughput and refurbishment cost efficiency compared to the no-policy scenario—all without increasing costs relative to the 4-depot baseline.

While more aggressive configurations, such as starting with 7 depots and applying policy scheme 7, refurbish many more satellites and deliver slightly better emissions, they cost substantially more (16%-36%). An initial configuration of 4 warehouses emerges as a reasonable balance between emissions goals and cost feasibility. This experiment illustrates the diminishing marginal returns in emission reduction and refurbishment increase that larger servicing infrastructures provide.

Both the grow-to-4 policy configuration from experiment 3.2 and the initial-4-with-cost-neutral-policy configuration demonstrate measurable improvements over the 4-depot laissez-faire approach. The choice between them depends on stakeholder priorities and whether the government can subsidize the cost difference or if sufficient secondary revenue opportunities exist. The 7.80% cost difference represents approximately \$350 million, indicating the necessary budget for subsidies or secondary revenue to close the business case for a 4-warehouse servicing infrastructure. Future work could explore dynamic policy application—for instance, transitioning from a growth-

oriented policy to a cost-neutral policy once the servicing infrastructure matures.

Figure 4.12 shows the emissions vs. total costs pareto frontier for all Experiment 3.3 configurations, identifying the previously highlighted (D:1, C:1, Rf:1, Up:0, RP:2, P:1, OUF:75k) and (D:4, C:1, Rf:1, Up:0, RP:2, P:1, OUF:50k) configurations as pareto-optimal points. The frontier denoted by the black curve shows the main frontier with all configuration points while the purple dotted frontier shows only no-policy configurations. This shows the modest shift downwards in the plot afforded by policy implementation, especially for 4 and 7 initial depots, which represents further NOx decrease for the same cost.

Furthermore, this plot illustrates the diminishing returns in NOX emission reductions as more warehouses are added compared to the increase in total costs. The nearly vertical line connecting the 0-warehouse baseline with the D1 configuration clearly depicts the cost-parity that a single depot+cost-neutral policy configuration has with the baseline. It is worthwhile to note that all pareto optimal points feature Rf=1 and Up=0, showing the value of flexible satellite upgrades and the preference for Earth-based refurbishing over space-based refurbishing.




Figure 4.12: Pareto Frontier: Total NOx Emissions (kg) vs. Total Costs (\$)




Figure 4.13: Pareto Frontier: Total Satellite Refurbishments vs. Total Costs (\$)

Similarly, Figure 4.13 illustrates an upward shift in the refurbishment vs. cost Pareto frontier when policy is included in the design space. Variations of Policy 1 dominate the majority of Pareto frontier. Just as it was in the emissions vs. cost Pareto frontier, $Rf=1$ and $Up=0$ are featured in nearly every Pareto optimal point.

Experimental Observation 3.3

The framework mechanics allows the user to compare policy schemes and identify scenario-dependent reward/penalty schemes that establish environmentally-beneficial infrastructures in a more economically feasible way than would occur through laissez-faire development alone. The policies don't eliminate costs but they reduce the economic strain of achieving better sustainability outcomes. Furthermore, the framework identifies the necessary dual-use project budget that would close the business case gap for larger infrastructures.

Table 4.26: Experiment 3.3: Total Cost

Top Performing Metric	Configuration	Percentile	Total Cost Value (\$)	Total Cost CI Lower	Total Cost CI Upper	Total Cost MoE%	Total Cost % vs Baseline	Total Cost p-value (Baseline)	Total Cost Sig? (Baseline)	Total Cost % vs Best No-Policy	Total Cost p-value (No-Policy)	Total Cost Sig? (No-Policy)
0-Warehouse Baseline	d:0 up:0 rf:0 rpo:0 P:0 OUF:0 Fine:0 Rebate:0	avg	4.50E+09	4.43E+09	4.58E+09	1.702	N/A	N/A	N/A	-0.322	0.780	No
		5th	4.02E+09	3.98E+09	4.14E+09	1.948	N/A	N/A	N/A	-2.534	0.780	No
		95th	5.11E+09	4.84E+09	5.12E+09	2.708	N/A	N/A	N/A	1.276	0.780	No
Best Cost No-Policy	d:1 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	4.52E+09	4.45E+09	4.59E+09	1.547	0.323	0.780	No	N/A	N/A	N/A
		5th	4.12E+09	4.05E+09	4.15E+09	1.213	2.600	0.780	No	N/A	N/A	N/A
		95th	5.05E+09	4.77E+09	5.13E+09	3.633	-1.260	0.780	No	N/A	N/A	N/A
Best for Total Cost	d:1 up:0 rf:1 rpo:2 P:1 OUF:7.5e4 R:1 Fine:0 Rebate:0	avg	4.85E+09	4.78E+09	4.93E+09	1.518	7.804	0.000	Yes	7.457	0.000	Yes
		5th	4.11E+09	4.03E+09	4.15E+09	1.426	2.351	0.887	No	-0.242	0.887	No
		95th	5.05E+09	4.78E+09	5.12E+09	3.395	-1.122	0.887	No	0.140	0.887	No
Best for Total Emissions	d:7 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	4.52E+09	4.45E+09	4.59E+09	1.542	0.332	0.774	No	0.009	0.994	No
		5th	4.80E+09	4.76E+09	4.88E+09	1.251	19.353	0.000	Yes	16.329	0.000	Yes
		95th	5.73E+09	5.57E+09	5.78E+09	1.825	12.047	0.000	Yes	13.477	0.000	Yes
Optimal for Cost/Refurbish	d:4 up:1 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	4.84E+09	4.77E+09	4.92E+09	1.507	7.618	0.000	Yes	7.272	0.000	Yes
		5th	5.22E+09	5.15E+09	5.29E+09	1.383	29.851	0.000	Yes	26.561	0.000	Yes
		95th	6.30E+09	6.08E+09	6.31E+09	1.825	23.307	0.000	Yes	24.881	0.000	Yes
Best for Sats Refurb (Space)	d:7 up:2 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	6.13E+09	6.04E+09	6.21E+09	1.358	36.116	0.000	Yes	35.678	0.000	Yes
		5th	5.62E+09	5.44E+09	5.67E+09	2.087	39.725	0.000	Yes	36.184	0.000	Yes
		95th	6.67E+09	6.56E+09	6.71E+09	1.124	30.588	0.000	Yes	32.255	0.000	Yes
Best for Num Sat Refurb and Cost/Collection	d:7 up:0 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	6.00E+09	5.92E+09	6.08E+09	1.314	33.271	0.000	Yes	32.842	0.000	Yes
		5th	5.53E+09	4.35E+09	4.45E+09	1.132	10.223	0.000	Yes	7.430	0.000	Yes
		95th	6.55E+09	5.15E+09	5.43E+09	2.621	5.262	0.000	Yes	6.606	0.000	Yes
Weighted Multi-Objective	d:4 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	4.85E+09	4.78E+09	4.93E+09	1.518	7.804	0.000	Yes	7.457	0.000	Yes
		5th	4.41E+09	4.36E+09	4.48E+09	1.361	9.634	0.000	Yes	6.857	0.000	Yes
		95th	5.35E+09	5.21E+09	5.47E+09	2.385	4.689	0.000	Yes	6.025	0.000	Yes
Weighted Multi-Obj (1-Depot Variant)	d:1 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	4.52E+09	4.45E+09	4.59E+09	1.542	0.332	0.774	No	0.009	0.994	No
		5th	4.13E+09	4.04E+09	4.15E+09	1.289	2.736	0.774	No	0.133	0.994	No
		95th	5.04E+09	4.78E+09	5.12E+09	3.387	1.397	0.774	No	-0.138	0.994	No
Weighted Multi-Obj (No-Policy Variant)	d:4 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	4.84E+09	4.77E+09	4.92E+09	1.507	7.618	0.000	Yes	7.272	0.000	Yes
		5th	4.43E+09	4.35E+09	4.45E+09	1.132	10.223	0.000	Yes	7.430	0.000	Yes
		95th	5.38E+09	5.15E+09	5.43E+09	2.621	5.262	0.000	Yes	6.606	0.000	Yes

Table 4.27: Experiment 3.3: Total NOx Emissions

Top Performing Metric	Configuration	Percentile	Total Emissions Value	Total Emissions CI Lower	Total Emissions CI Upper	Total Emissions MoE%	Total Emissions % vs Baseline	Total Emissions p-value (Baseline)	Total Emissions Sig? (Baseline)	Total Emissions % vs Best No-Policy	Total Emissions p-value (No-Policy)	Total Emissions Sig? (No-Policy)
0-Warehouse Baseline	d:0 up:0 rf:0 rpo:0 P:0 OUF:0 Fine:0 Rebate:0	avg	3.49E+06	3.42E+06	3.56E+06	2.004	N/A	N/A	N/A	7.587	0.000	Yes
		5th	3.01E+06	2.98E+06	3.12E+06	2.237	N/A	N/A	N/A	17.380	0.000	Yes
		95th	3.94E+06	3.83E+06	4.01E+06	2.275	N/A	N/A	N/A	5.170	0.000	Yes
Best Cost No-Policy	d:1 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	3.24E+06	3.16E+06	3.32E+06	2.556	-7.052	0.000	Yes	N/A	N/A	N/A
		5th	2.57E+06	2.49E+06	2.84E+06	6.683	-14.806	0.000	Yes	N/A	N/A	N/A
		95th	3.75E+06	3.64E+06	3.84E+06	2.671	-4.916	0.000	Yes	N/A	N/A	N/A
Best for Total Cost	d:1 up:0 rf:1 rpo:2 P:1 OUF:7.5e4 R:1 Fine:0 Rebate:0	avg	3.14E+06	3.06E+06	3.21E+06	2.355	-9.939	0.000	Yes	-3.106	0.073	No
		5th	2.64E+06	2.52E+06	2.88E+06	6.788	-12.475	0.000	Yes	2.736	0.893	No
		95th	3.72E+06	3.65E+06	3.87E+06	2.938	-5.592	0.000	Yes	-0.711	0.893	No
Best for Total Emissions	d:7 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	3.25E+06	3.17E+06	3.32E+06	2.424	-6.858	0.000	Yes	0.208	0.906	No
		5th	2.63E+06	2.43E+06	2.74E+06	5.785	-12.788	0.000	Yes	2.369	0.000	Yes
		95th	3.40E+06	3.36E+06	3.46E+06	1.452	-13.832	0.000	Yes	-9.377	0.000	Yes
Optimal for Cost/Refurbish	d:4 up:1 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	3.18E+06	3.10E+06	3.25E+06	2.311	-8.814	0.000	Yes	-1.896	0.270	No
		5th	2.79E+06	2.56E+06	2.87E+06	5.525	-7.578	0.000	Yes	8.485	0.826	No
		95th	3.59E+06	3.52E+06	3.80E+06	3.855	-8.990	0.000	Yes	-4.285	0.826	No
Best for Sats Refurb (Space)	d:7 up:2 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	3.22E+06	3.16E+06	3.27E+06	1.719	-7.678	0.000	Yes	-0.673	0.663	No
		5th	2.85E+06	2.76E+06	2.90E+06	2.408	-5.487	0.000	Yes	10.939	0.663	No
		95th	3.54E+06	3.47E+06	3.67E+06	2.853	-10.242	0.000	Yes	-5.601	0.663	No
Best for Num Sat Refurb and Cost/Collection	d:7 up:0 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	3.09E+06	3.04E+06	3.15E+06	1.844	-11.266	0.000	Yes	-4.534	0.004	Yes
		5th	2.68E+06	2.49E+06	2.75E+06	5.029	-13.947	0.000	Yes	1.009	0.270	No
		95th	3.48E+06	3.52E+06	3.64E+06	1.582	-8.603	0.000	Yes	-3.878	0.270	No
Weighted Multi-Objective	d:4 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	3.14E+06	3.06E+06	3.21E+06	2.355	-9.939	0.000	Yes	-3.106	0.073	No
		5th	2.66E+06	2.45E+06	2.78E+06	6.265	-11.579	0.000	Yes	3.788	0.073	No
		95th	3.60E+06	3.51E+06	3.67E+06	2.263	-8.803	0.000	Yes	-4.088	0.073	No
Weighted Multi-Obj (1-Depot Variant)	d:1 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	3.25E+06	3.17E+06	3.32E+06	2.424	-6.858	0.000	Yes	0.208	0.906	No
		5th	2.61E+06	2.46E+06	2.82E+06	6.795	-13.371	0.000	Yes	1.685	0.906	No
		95th	3.72E+06	3.61E+06	3.78E+06	2.232	-5.781	0.000	Yes	-0.910	0.906	No
Weighted Multi-Obj (No-Policy Variant)	d:4 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	3.18E+06	3.10E+06	3.25E+06	2.311	-8.814	0.000	Yes	-1.896	0.270	No
		5th	2.59E+06	2.49E+06	2.75E+06	5.029	-13.947	0.000	Yes	1.009	0.270	No
		95th	3.61E+06	3.52E+06	3.64E+06	1.582	-8.603	0.000	Yes	-3.878	0.270	No

Table 4.28: Experiment 3.3: Number of Refurbishments

Top Performing Metric	Configuration	Percentile	Num Sat Refurb Value	Num Sat Refurb CI Lower	Num Sat Refurb CI Upper	Num Sat Refurb MoE%	Num Sat Refurb % vs Best No-Policy	Num Sat Refurb p-value (No-Policy)	Num Sat Refurb Sig? (No-Policy)
0-Warehouse Baseline	d:0 up:0 rf:0 rpo:0 P:0 OUF:0 Fine:0 Rebate:0	avg	0.00E+00	0.00E+00	0.00E+00	N/A	N/A	N/A	N/A
		5th	0.00E+00	0.00E+00	0.00E+00	N/A	N/A	N/A	N/A
		95th	0.00E+00	0.00E+00	0.00E+00	N/A	N/A	N/A	N/A
Best Cost No-Policy	d:1 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	5.75E+01	4.75E+01	6.75E+01	17.402	N/A	N/A	N/A
		5th	0.00E+00	0.00E+00	1.69E+01	-	N/A	N/A	N/A
		95th	1.45E+02	9.93E+01	1.61E+02	21.327	N/A	N/A	N/A
Best for Total Cost	d:1 up:0 rf:1 rpo:2 P:1 OUF:7.5e4 R:1 Fine:0 Rebate:0	avg	2.62E+02	2.46E+02	2.78E+02	6.218	355.752	0.000	Yes
		5th	1.00E+01	9.00E+00	3.53E+01	131.250	inf	0.011	Yes
		95th	1.55E+02	1.31E+02	1.72E+02	13.306	7.155	0.011	Yes
Best for Total Emissions	d:7 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	7.16E+01	6.22E+01	8.11E+01	13.197	24.596	0.042	Yes
		5th	2.13E+02	1.19E+02	2.68E+02	35.036	inf	0.000	Yes
		95th	4.57E+02	4.32E+02	4.75E+02	4.717	216.212	0.000	Yes
Optimal for Cost/Refurbish	d:4 up:1 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	1.94E+02	1.76E+02	2.13E+02	9.567	238.186	0.000	Yes
		5th	2.63E+02	1.75E+02	3.12E+02	26.138	inf	0.000	Yes
		95th	5.25E+02	4.93E+02	5.76E+02	7.883	263.083	0.000	Yes
Best for Sats Refurb (Space)	d:7 up:2 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	3.79E+02	3.59E+02	3.98E+02	5.139	559.081	0.000	Yes
		5th	2.35E+02	1.61E+02	3.00E+02	29.556	inf	0.000	Yes
		95th	4.83E+02	4.61E+02	5.61E+02	10.302	234.186	0.000	Yes
Best for Num Sat Refurb and Cost/Collection	d:7 up:0 rf:1 rpo:2 P:7 OUF:2.5e4 Fine:10M Rebate:250k	avg	4.12E+02	3.91E+02	4.33E+02	5.090	616.124	0.000	Yes
		5th	2.56E+02	3.50E+01	1.00E+02	38.770	inf	0.000	Yes
		95th	5.28E+02	2.80E+02	3.36E+02	9.167	112.859	0.000	Yes
Weighted Multi-Objective	d:4 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	2.62E+02	2.46E+02	2.78E+02	6.218	355.752	0.000	Yes
		5th	1.60E+02	1.40E+02	1.81E+02	12.789	inf	0.000	Yes
		95th	3.81E+02	3.52E+02	3.87E+02	4.664	163.083	0.000	Yes
Weighted Multi-Obj (1-Depot Variant)	d:1 up:0 rf:1 rpo:2 P:1 OUF:5e4 R:1 Fine:0 Rebate:0	avg	7.16E+01	6.22E+01	8.11E+01	13.197	24.596	0.042	Yes
		5th	1.00E+01	0.00E+00	3.58E+01	179.000	inf	0.042	Yes
		95th	1.48E+02	1.24E+02	1.61E+02	12.606	2.281	0.042	Yes
Weighted Multi-Obj (No-Policy Variant)	d:4 up:0 rf:1 rpo:2 P:0 OUF:0 Fine:0 Rebate:0	avg	1.94E+02	1.76E+02	2.13E+02	9.567	238.186	0.000	Yes
		5th	8.39E+01	3.50E+01	1.00E+02	38.770	inf	0.000	Yes
		95th	3.08E+02	2.80E+02	3.36E+02	9.167	112.859	0.000	Yes

4.4.4 Experimental Support for Hypothesis 3

This section examines the extent to which experimental findings from Experiments 3.1, 3.2, and 3.3 substantiate Hypothesis 3, which posits that market-based reward/penalty schemes can establish economically feasible OOS infrastructures that yield better sustainability metrics than laissez-faire approaches. The evaluation is structured around four dimensions: the existence of effective policy schemes, their economic feasibility, their environmental superiority relative to laissez-faire alternatives, and the conditions of their performance.

Existence of Effective Policy Schemes

Experiment 3.1 demonstrated that the parametric reward/penalty framework successfully models diverse policy mechanisms with measurably different impacts on system performance. The down-selection analysis identified 18 policy variants across eight distinct schemes, each exhibiting characteristic patterns of statistical significance (based on ANOVA-like variance analysis) for different performance metrics. Policy 1 (OUF with refund conditions) achieved significance for emissions across all cost and emissions percentiles. Policies 3, 6, and 7 (fine-based mechanisms) showed significance at extreme percentiles, indicating their role in managing worst-case compliance scenarios. Policy 5 (tax-based mechanism) exhibited rebate significance across all cost percentiles. These varied significance patterns demonstrate that different policy mechanisms target different aspects of system performance.

Cost-Neutral Pathways to Enhanced Performance

Experimental Support 3.1 provides direct evidence for hypothesis validation through the \$75,000 OUF scheme with refund conditional on satellite collection (Policy 1.2). This configuration achieved cost-neutrality while delivering meaningful operational improvements in refurbishment throughput. However, the hypothesis qualification of

“better sustainability metrics than laissez-faire” requires careful interpretation. Policy 1.2 does not deliver statistically measurable improvements in average emissions compared to the best-performing no-policy configuration given the number of scenarios tested. The environmental benefits manifest primarily through enhanced refurbishment throughput (29.47% increase over the top no-policy configuration, $p = 0.038$) rather than through shifts in average emissions performance tendencies.

Strategic Pairing of Policy with Infrastructure Flexibility

Experimental Support 3.2 reveals that policy-enabled flexible warehouse roll-out allows the framework to reach a target infrastructure size at lower total cost compared to initializing with equivalent infrastructure at the simulation start. Policy 2.2 (grow-to-7) achieved 7.5 average warehouses with an average total cost 11.47% higher than the 0-warehouse baseline, compared to the no-policy configuration that starts with 7 warehouses which is 16.22% more expensive than the 0-warehouse baseline. Similarly, Policy 3.1 (grow-to-4) created 4 warehouses on average with an average total cost of 5.29% greater than the 0-warehouse baseline, compared to the no-policy configuration that starts with 4 warehouses with a 7.59% greater average total cost.

However, this cost-efficient incremental deployment approach exhibits losses in emission reductions and refurbishment capabilities relative to configurations initialized with equivalent infrastructure capacity. For instance, the 4-depot no-policy configuration achieves -10.05% emissions reduction compared to the 0-warehouse baseline and delivers 198 average refurbishments while the grow-to-4 configuration has a -6.02% emissions reduction and 120 average refurbishments. To achieve substantial structural change in environmental metrics, a greater number of warehouses must be present at the start of the simulation.

This represents a fundamental trade-off: cost-efficient incremental deployment sacrifices near-term environmental benefits that could be achieved through larger ini-

tial infrastructure investments. The finding partially contradicts the hypothesis by demonstrating that not all reward/penalty schemes yield “better sustainability metrics” than alternatives. The definition of “better” depends critically on whether one prioritizes cost efficiency or immediate environmental performance. This framework provides a testbed to uncover these trade-offs.

Policy-Enabled Initial Infrastructure Investment

Experimental Support 3.3 provides the strongest validation of Hypothesis 3 by demonstrating how policy can enable larger initial infrastructure investments while maintaining economic feasibility. The configuration with 4 initial depots and Policy 1 (D:4, CAAS:1, Ref:1, Upg:0, P:1, RPO:2, OUF:5e4) achieves -9.94% average emissions reduction relative to the 0-warehouse baseline [95% CI: [3.06M, 3.21M], %MoE: 2.36], while maintaining +7.80% cost difference [95% CI: [4.78B, 4.93B], %MoE: 1.52] (\$350M over 30 years). This configuration outperforms both the 1-depot policy variant and the 4-depot no-policy variant in sustainability metrics.

The 4-depot policy configuration delivers 262 average total refurbishments compared to 194 for the 4-depot no-policy equivalent, while achieving cost per refurbishment with -95.79% difference compared to the top cost no-policy configuration (versus -79.45% for the 4-depot no-policy variant). The 7.80% cost difference represents the necessary budget for subsidies or secondary revenue opportunities to close the business case gap for a 4-warehouse servicing infrastructure. The framework identifies this business case gap, enabling policymakers to design interventions that address the specific economic barriers to infrastructure development.

The 1-depot variant of this policy (D:1, CAAS:1, Ref:1, Upg:0, P:1, RPO:2, OUF:5e4) costs about the same as the 0-warehouse baseline while achieving -6.86% emissions reduction and refurbishing 72 satellites on average. This demonstrates sustainability improvements can be achieved at cost-neutrality, though the 4-depot

configuration provides far greater absolute benefits.

Priority Dependence and Stakeholder Alignment

The experimental results demonstrate clear priority dependence in policy performance. Cost-neutral schemes align with stakeholders prioritizing minimal economic disruption and increased servicing throughput. More ambitious schemes align with stakeholders willing to accept higher costs for enhanced environmental outcomes or infrastructure development. The down-selection of multiple configurations per policy (x.1, x.2, x.3) enables systematic exploration of this trade-off space. Configuration x.1 variants consistently prioritize cost-weighted balance, x.2 variants emphasize environmental metrics at acceptable cost premiums, and x.3 variants explore specialized objectives such as increasing refurbishment throughput.

4.4.5 Hypothesis 3 Substantiation

Based on experimental evidence, Hypothesis 3 is validated with important qualifications:

1. **Existence of Effective Policies:** The framework successfully models parametric reward/penalty schemes with measurably different impacts on infrastructure feasibility and sustainability metrics.
2. **Economic Feasibility:** Cost-neutral policies (Policy 1 with \$75k OUF and refund) achieve near cost-parity with baselines while improving refurbishment throughput by 29.47%. However, policies that enable substantially better environmental outcomes require \$350M+ subsidies or secondary revenue over 30 years (7.80% cost increase for the 4-depot Policy 1 configuration).
3. **Environmental Superiority:** Policy schemes demonstrate better sustainability metrics than laissez-faire alternatives primarily when coupled with larger

initial infrastructure investments rather than through cost-neutral incremental deployment alone. Cost-neutral policies improve refurbishment throughput but do not provide statistically measurable improvements in average emissions given the tested scenario count.

4. **Priority Dependence:** Policy effectiveness varies systematically with infrastructure scale and stakeholder priorities.

Addressing Research Question 4: Policy Design Principles

The experimental results provide direct insights for Research Question 4 regarding optimal combinations and calibrations of policy parameters:

For stakeholders prioritizing minimal cost impact: Policy 1.2 (\$75k OUF with refund upon collection) provides the optimal approach, achieving cost-neutrality while improving refurbishment throughput by 29.47% ($p = 0.038$) and decreasing risk at extreme percentiles.

For stakeholders prioritizing environmental outcomes within moderate cost constraints: Policy 1 (\$50k OUF with refund) applied to 4-depot initial infrastructure provides the best balance, achieving -9.94% decrease in average emissions [95% CI: [3.06M, 3.21M], %MoE: 2.36] at a 7.80% cost increase [95% CI: [4.78B, 4.93B], %MoE: 1.52] (\$350M over 30 years). This configuration delivers 262 average refurbishments with -95.79% cost per refurbishment difference compared to the top no-policy configuration.

For stakeholders prioritizing infrastructure development: Policies 2.2 and 3.1 enable flexible warehouse roll-out, reaching target warehouse numbers at lower total cost than initialization with equivalent capacity. Policy 2.2 grows to 7.5 warehouses at 11.47% cost increase (versus 16.22% for initialized 7-depot no-policy). Policy 3.1 grows to 4 warehouses at 5.29% cost increase (versus 7.59% for initialized 4-depot no-policy). These policies distribute infrastructure investment across time while building

operational capacity, though at the expense of near-term environmental benefits.

The experimental results demonstrate that no single policy configuration simultaneously optimizes all objectives, but that the parametric framework enables systematic identification of configurations aligned with specific stakeholder priorities. This addresses Research Question 4 by showing that high-performing policy design depends on stakeholder priorities regarding cost-neutrality versus environmental ambition, and the framework provides decision-makers with quantified trade-offs to support policy selection. The Pareto frontier analysis in Experiment 3.3 further illustrates these trade-offs, showing the diminishing returns in NOx emission reductions as more warehouses are added compared to the increase in total costs, with all Pareto-optimal points featuring flexible satellite upgrades ($Rf=1$) and no warehouse upgrades ($Upg=0$).

4.5 Experimental Support for Overarching Hypothesis

Recall from chapter 2 the overarching observation and the subsequent overarching research question:

Overarching Observation

The process of motivating and establishing circular space economies in LEO is a complex system of systems problem that requires analysis from the technical, financial, and policy perspectives. This thesis aims to provide a systems-level screening framework that evaluates the interaction between novel OOS CONOPs, flexible options, and various policy schemes in order to path-find strategies and infrastructures that could improve the case for OOS in LEO.

Overarching Research Question

Which flexible option, or set of options, consistently improves the economic value and environmental impact of LEO-based OOS over a range of potential policies and future scenarios sampled from multi-domain uncertainty?

Having conducted experiments 1 through 3 and evaluated their degree of support for their respective hypotheses, we now revisit the Overarching Hypothesis:

Overarching Hypothesis

If a flexibility framework for LEO-based OOS incorporates multiple uncertain variables, policy impact, novel design concepts like collection hubs, and allows for multiple combinatorial options, then there will exist an option or set of options that provides a viable and sustainable private infrastructure for a circular space economy in LEO.

The experimental results from Hypotheses 1 through 3 provide comprehensive support for the overarching hypothesis that a flexibility framework incorporating multiple uncertain variables, policy impact, and combinatorial options can identify viable pathways toward sustainable infrastructure for circular space economies in LEO. This support manifests through several findings that challenge prevailing assumptions in the literature, reveal non-obvious system interactions, and quantify specific technical and policy mechanisms that enable both economic feasibility and environmental sustainability. The major findings are presented below, organized by experiment.

4.5.1 Major Findings from Experiment 1

Major Finding 1.1: Earth-Based Refurbishment Dominates Space-Based Servicing

Collection for Earth-based refurbishment paired with in-space ADR vehicle servicing demonstrate greater environmental and operational value than in-space satellite servicing. Warehouse upgrades enabling in-space satellite servicing show limited sta-

statistical significance for emissions reduction when other parameters are held constant, while Earth-return dual-mission operations emerge as the primary emissions reduction mechanism. This counterintuitive result reflects fundamental economic realities: declining satellite manufacturing costs combined with declining launch costs make in-space component replacement economically challenging compared to Earth-based refurbishment that leverages existing terrestrial infrastructure. The critical enabling technology is reusable second-stage launch vehicles capable of transporting satellites to Earth intact, which reduces atmospheric reentry emissions while enabling batch processing of collected satellites. This finding redirects technology development priorities from complex on-orbit satellite servicing capabilities toward reusable launch vehicle architectures and orbital warehouses.

Major Finding 1.2: RPO Capability Prioritization Over Satellite Refuelability

Rendezvous and proximity operations (RPO) capability available from mission inception consistently demonstrates higher value than satellite refuelability across all experiments. This result is surprising given that existing literature and industry attention emphasizes making satellites serviceable rather than making them RPO-capable for servicing operations. All top-performing configurations across experiments feature RPO=2 (immediately available), identifying the capability as a flexibility enabler that delivers both economic and sustainability benefits. RPO capability improves operational efficiency throughout the CAAS ecosystem by reducing ADR vehicle downtime, enabling satellites to maneuver to collection locations, facilitating more efficient warehouse operations, and ultimately improving refurbishment throughput.

Major Finding 1.3: Single Depot Captures Majority of System Benefits

A single orbital warehouse captures much of the operational efficiency and environmental benefits of multi-depot architectures while avoiding multiplicative capital

costs, despite intuition suggesting that distributed infrastructure would be necessary for mega-constellation servicing. The cost-optimal CAAS configuration (D1, C1, Rf0, Up0, RP2) achieves a modest -6.48% emissions reduction ($p < 0.001$) and enables an average of 65 satellite refurbishments over 30 years at cost parity with the 0-warehouse baseline (+0.76%, $p = 0.74$ not statistically significant). Scaling to 4 depots improves emissions and refurbishments but increases costs, while 7 depots cost 16%-36% more than baseline while providing only modest additional emissions benefits compared to 4 depots. This pattern of diminishing marginal returns establishes that initial CAAS deployments should prioritize single-depot architectures to maximize cost-effectiveness, with expansion to multiple depots reserved for scenarios where substantial government subsidy or secondary revenue streams justify the multiplicative infrastructure investment. These findings suggest that not all satellites in a constellation need to be collected for a CAAS infrastructure to be worthwhile.

Major Finding 1.4: Partial Constellation Servicing Provides Measurable Benefits

Servicing a relatively small portion of a mega-constellation's satellites over the 30-year analysis period provides statistically significant and measurable operational improvements. The cost-optimal single-depot configuration refurbishes an average of 65 satellites with 95% CI [54, 75] from an 18-plane, 648-satellite constellation while achieving -6.48% emissions reduction ($p < 0.001$). This finding is significant because it lowers the threshold for viable CAAS business models; comprehensive constellation-wide servicing is not a prerequisite for demonstrating value. Early-stage systems can target high-value satellites or technology demonstration opportunities to achieve measurable returns while building operational experience and validating technical capabilities. As technologies mature and costs decline, future constellations can be intentionally designed for CAAS compatibility, creating a pathway for incremental adoption rather than requiring constellation redesign and reconfiguration as a precondition for CAAS

implementation.

4.5.2 Major Findings from Experiment 2

Major Finding 2.1: Flexible Satellite Serviceability Outperforms Fixed Deployment

Flexible satellite serviceability deployment ($Rf=1$) consistently outranks both immediate deployment ($Rf=2$) and never deploying the capability ($Rf=0$) across 80 uncertainty scenarios, though the absolute cost difference (0.243 percentage points improvement over inflexible CAAS) is too small to quantify with statistical significance given the sample size. Ranking convergence analysis confirms stable second-place performance for the flexible configuration ($D1, C1, Rf1, Up0, RP2$) compared to third-place for the inflexible equivalent ($D1, C1, Rf0, Up0, RP2$). This finding demonstrates the value of adaptive deployment strategies that respond to evolving technology costs and operational experience rather than committing to fixed architectures at mission inception. This finding, as well as the finding that the CAAS CONOPs is strengthened by its ADR-deployment capabilities, improved by its use of reusable second stage vehicles, and threatened by lower failure rates, demonstrates this framework's ability to identify top-performing aspects of the CAAS CONOPs, their relative importance, and their conditions.

4.5.3 Major Findings from Experiment 3

Major Finding 3.1: Cost-Neutral Policies Improve Refurbishment Throughput

The framework identifies specific policy parameter combinations that achieve cost neutrality over 30-year horizons while significantly improving refurbishment throughput, though without providing statistically measurable emissions benefits beyond no-policy CAAS configurations. Policy 1.2 (\$75k annual Orbital Use Fee with collection rebate) achieves cost parity with the 0-warehouse baseline while delivering 29.47%

increase in satellite refurbishments ($p = 0.038$, 95% CI: [65, 93], MoE: 16.9%) compared to the top no-policy configuration. However, emissions improvements remain too small to measure with statistical significance given 80 scenarios ($p = 0.56$). This reveals a fundamental bifurcation in policy mechanism design: cost-neutral hybrid fee-rebate structures enhance refurbishment throughput without shifting central environmental performance tendencies, while emissions-focused policies require different architectures involving substantial upfront infrastructure investment or more aggressive enforcement parameters. The existence of cost-neutral pathways demonstrates that catalyzing sustainable space infrastructure development is possible without imposing permanent fiscal burden on government or industry, provided stakeholders prioritize refurbishment throughput over aggressive emission reduction targets. Pairing policy 1 with a 4-depot initial infrastructure provides the emissions benefits of a larger CAAS infrastructure while improving throughput and managing costs relative to a laissez-faire 4-depot configuration. Overall, the framework provides a testbed to help policymakers understand the tradeoffs of varying configurations and policy parameters.

4.5.4 Minor Supporting Findings

The following findings provide additional context but represent more intuitive results or validate expected patterns:

- **Satellite failure rate assumptions and space-based ADR deployment appear to influence CAAS competitiveness.** While more scenarios are required for better confidence in the absolute cost differences, the sensitivity tests suggest that CAAS cost disadvantage increases as the initial failure rate decreases, indicating that the CAAS value proposition depends fundamentally on failure frequency. Additionally, when ADR vehicles launch exclusively from Earth rather than deploying from orbital warehouses, the cost difference

increases relative to baseline, indicating that orbital ADR staging capability improves cost competitiveness.

- **Decreasing ADR Vehicle Cost and Improving Earth-Refurbishment Streamlining Increase CAAS Competitiveness.** The total cost rank order placed the CAAS configuration above the 0-warehouse baseline in sensitivity tests that decreased the minimum achievable ADR vehicle cost (via the learning curve effect) and reduced Earth-refurbishment costs for refuelable and repairable satellites. This provides guidance for improving ADR vehicle manufacturing costs and investing in streamlined Earth-refurbishment capabilities for returned satellites.
- **Deterministic cost perturbations produce minimal ranking disruption.** The cost-superior configuration (D1, Rf1, Up0, RP2) maintains stability across all perturbation tests, with cost differences remaining within a small percentage of the 0-warehouse baseline.
- **Policy-enabled flexible warehouse deployment trades cost efficiency for near-term environmental impact.** Growing to 4 warehouses via Policy 3.1 costs 5.29% more than baseline versus 7.59% for initializing with 4 warehouses, but achieves -6.02% emissions reduction compared to -8.81% for initialized configuration and delivers 120 versus 194 average refurbishments.
- **Multi-echelon sparing without CAAS provides no value.** The best no-CAAS depot configuration incurs a statistically significant 3.29% cost penalty ($p = 0.025$) while achieving limited emissions difference ($p = 0.89$, not significant), validating that passive spare depot infrastructure is insufficient compared to the CAAS concept for the given use case.

4.5.5 Integration: Addressing the Overarching Research Question

The original problem statement identified the need for systems-level analysis to evaluate interactions between novel OOS concepts, flexible options, and policy schemes to pathfind strategies and infrastructures that could improve the case for OOS in LEO. The overarching research question asked: *Which flexible option, or set of options, consistently improves the economic value and environmental impact of LEO-based OOS over a range of potential policies and future scenarios sampled from multi-domain uncertainty?*

The experimental evidence provides a definitive answer through systematic identification of technical configurations, deployment strategies, and policy mechanisms that enable viable circular space economy pathways. Single-depot CAAS with active collection, space-based ADR hosting, flexible satellite serviceability ($R_f=1$), pre-initialized RPO capability ($R_P=2$), and Earth-return refurbishment focus consistently emerges as the cost-superior configuration. This architecture achieves -6.48% emissions reduction ($p < 0.001$) at cost parity with baseline (+0.76%, $p = 0.74$ not significant) while enabling 65 average satellite refurbishments [95% CI: 54-75] over 30 years. Including policy 1.2 (\$75k annual Orbital Use Fee with collection rebate) achieves cost parity with the 0-warehouse baseline while delivering 29.47% increase in satellite refurbishments.

The framework identifies two distinct policy archetypes aligned with different stakeholder priorities. Cost-neutral mechanisms (\$75k OUF with collection rebate) improve operational metrics without fiscal burden, achieving 29.47% refurbishment throughput increase ($p = 0.038$) while maintaining cost parity. Emissions-focused mechanisms (Policy 1 + 4-Depot Infrastructure) require \$350M business case gap closure but enable -9.94% emissions reduction and 262 satellite refurbishments through larger infrastructure initialization. Growth-focused mechanisms reduce the relative cost burden of initialized large infrastructures from the start, but provide less sus-

tainability benefits.

The analysis reveals several counterintuitive system interactions that challenge conventional assumptions. Earth-based refurbishment dominates space-based servicing for both emissions and cost efficiency, while partial constellation collection delivers measurable sustainability improvements without requiring comprehensive fleet-wide implementation, offering a first step towards improved circularity for the OneWeb constellation as-is.

From a methodological perspective, ranking convergence analysis enables confident technology selection when statistical power limitations prevent significance testing of small absolute differences, providing practical decision-making frameworks for high-complexity system-of-systems problems where thousands of Monte Carlo samples may be computationally prohibitive.

4.5.6 Implications for Circular Space Economy Development and Closing Remarks

The experimental results demonstrate that viable pathways to sustainable LEO infrastructure exist, but their success depends on precise alignment of technical architecture, deployment strategy, policy environment, and stakeholder priorities:

- **For cost-competitive sustainability improvement with minimal risk:**
Deploy single-depot CAAS with flexible satellite serviceability, pre-initialized RPO, and Earth-return refurbishment focus, potentially augmented with \$75k OUF cost-neutral policy to enhance refurbishment throughput.
- **For ambitious environmental outcomes accepting moderate cost increases:** Initialize 4-depot infrastructure with \$50k OUF policy, recognizing the \$350M business case gap that must be closed through government subsidy or secondary revenue streams.
- **For long-term infrastructure development prioritizing cost efficiency:**

Deploy policy-enabled flexible warehouse growth, accepting reduced near-term environmental benefits in exchange for time-value-of-money advantages from delayed capital expenditure.

The framework does not identify a single universally-optimal solution, nor does it eliminate the fundamental trade-offs between cost, environmental impact, and operational capability that characterize complex system-of-systems problems. Rather, it provides the analytical machinery to quantify these trade-offs with statistical rigor, identify non-obvious mechanisms and dependencies that contradict prevailing literature assumptions, and map performance across multi-domain uncertainty distributions that represent plausible futures for space operations.

Two overarching insights emerge that directly address the original problem statement. First, circular space economy infrastructure can achieve cost competitiveness with traditional make-use-dispose approaches when properly architected around Earth-return refurbishment, space-based ADR hosting, and selective flexible deployment rather than comprehensive in-space servicing capabilities emphasized in existing literature.

Second, viable pathways exist across the stakeholder priority spectrum from risk-averse cost-neutral operational enhancement to aggressive emissions reduction, with the framework providing decision-makers the quantified trade-offs necessary to select strategies aligned with their specific political, economic, and environmental contexts.

The overarching hypothesis is thereby substantiated: flexibility frameworks incorporating combinatorial technical options, multi-domain uncertainty, flexible deployment strategies, and parametric policy mechanisms can identify viable pathways toward circular space economies. “Viability” manifests not as universal superiority across all metrics but as context-dependent alignment between stakeholder priorities, technical configurations, deployment strategies, and policy environments. The experimental evidence demonstrates that the framework successfully performs its intended

screening function: illuminating the multi-dimensional trade-space where informed strategic choices can be made, revealing counterintuitive system behaviors that challenge prevailing assumptions, and quantifying the specific technical capabilities and policy interventions required to close identified business case gaps.

This systems-level screening capability, validated through three experimental hypotheses spanning technical architecture evaluation (Hypothesis 1), flexibility enabler/mechanism identification (Hypothesis 2), and policy-flexibility interaction analysis (Hypothesis 3), confirms that integrated frameworks incorporating the complexity of LEO-based OOS systems are not merely useful but necessary tools for pathfinding sustainable space operations. The alternative of optimizing individual subsystems or policy mechanisms in isolation fails to capture the critical interactions and non-obvious dependencies that ultimately determine system-level viability in realistic operational environments characterized by technological uncertainty, cost dynamics, and competing stakeholder objectives.

CHAPTER 5

CONCLUSION

The space industry is consistently at the forefront of technology, inspiring generations and improving human knowledge. The very industry that puts people on the moon and gazes into the depths of space could also be the industry that creates the next circular economy. Such a system would not only reduce pollution, but also expand our capacity for future space exploration as we improve our resource utilization in a resource-starved environment. Space debris is a known problem and atmospheric pollution is a growing problem; improving the case for OOS in LEO could mitigate both of them. This thesis considers the three-pronged approach for improving OOS for LEO: better CONOPs, better strategies, and better policies. Underscoring the value of flexibility, the opportunity in uncertainty, and the complex web of interactions between systems and decision-makers, this thesis screens options that promote long-term sustainability of our activities beyond Earth.

The overarching research question driving this work asked: *Which flexible option, or set of options, consistently improves the economic feasibility and environmental impact of LEO-based OOS over a range of potential policies and future scenarios sampled from multi-domain uncertainty?* Through systematic exploration of Collection-as-a-Service (CAAS) concepts, flexibility mechanisms, and policy interventions using the Space JAWA framework, this thesis identifies viable pathways toward circular space economies while acknowledging the inherent trade-offs and context dependencies that shape infrastructure decisions.

5.1 Major Contributions and Key Findings

This research makes three major contributions:

First, it provides a comprehensive flexibility framework for LEO-based on-orbit servicing. This framework captures multiple uncertainty sources, such as launch costs and satellite manufacturing costs, technology obsolescence, and random failures, using Monte Carlo analysis across various uncertain scenarios. It models customer demand for OOS based on the state of uncertainty and interactions between an evolving satellite constellation and OOS infrastructure, including decision rules for incremental infrastructure deployment and satellite upgrade decisions. The framework enables comparison of flexible strategies to identify consistent value drivers across diverse futures.

Second, it contributes novel OOS elements and concepts of operations specifically tailored for LEO. This includes collection hubs with upgradable capabilities, cooperative maneuvering leveraging J2 perturbations for natural plane drift, temporary abandonment logistics, and Earth-return using reusable vehicles to eliminate atmospheric emissions. These concepts represent fundamental departures from existing GEO-focused OOS approaches, recognizing that LEO's unique characteristics require purpose-built solutions rather than adaptations of GEO practices.

Third, it provides a policy impact assessment methodology that evaluates market-based mechanisms, such as fees, subsidies, taxes, revenue-neutral designs, and their effects on economic viability and environmental performance. Together, these form a comprehensive framework for screening CONOPs, strategies, and policies that aim to promote sustainable space infrastructures under deep uncertainty. The framework explicitly quantifies trade-offs between economic and environmental objectives, identifies conditions under which different strategies excel, and illuminates policy designs that overcome market coordination failures.

5.1.1 Non-Intuitive Findings

Analysis across all experiments revealed several counter-intuitive findings that challenge prevailing assumptions or positions in space sustainability literature:

Collection for Earth-based refurbishment shows more potential than space-based refurbishment. This is surprising because much of the literature about creating circular space economies focuses on in-space manufacturing and servicing rather than Earth-return and refurbishment. However, declining satellite manufacturing costs combined with declining launch costs make in-space refurbishment economically challenging. Earth-based refurbishment leverages the decline in launch costs as well as existing terrestrial infrastructure. The critical enabler is reusable second-stage vehicles capable of transporting satellites to Earth intact.

RPO capability prioritization emerges as more valuable than satellite refuelability. This result is surprising because there is more literature and attention to making satellites serviceable than making them RPO-capable for the purposes of satellite servicing and improving circularity. Yet across all experiments and sensitivity analyses, rendezvous and proximity operations capability that is available from the start consistently showed better value than satellite refuelability. RPO capability improves operational efficiency throughout the CAAS ecosystem: it reduces ADR vehicle down-time and facilitates better refurbishment throughput.

Flexible refueling initialization proves more valuable than fixed deployment. While the cost difference was too small to quantify definitively with the number of Monte Carlo samples, stable and consistent rank-order analysis shows that flexible refueling provides a cost benefit over triggering the option from the start. The flexibility to defer refueling capability until its value justifies investment consistently outperforms either immediate deployment or never deploying the capability.

Another finding is that warehouse satellite servicing capability is not the dominant driver for emissions reduction. The experiments revealed that a dominant driver of

emissions reduction is the Earth-return vehicle technology maturation timeline, not the ability to refurbish satellites in the warehouses. A basic warehouse aggregating satellites for batch Earth-return provides nearly the same emissions benefits of a fully-capable servicing depot. This has important implications: early-stage systems should focus on satellite collection and ADR vehicle in-space servicing rather than sophisticated in-space servicing for satellites, which marks a shift from circular space economy literature that focuses on satellite servicing in space.

A one-warehouse CAAS infrastructure doesn't come close to collecting all satellites in a constellation, but even collecting a relatively small portion of a constellation's satellites over the analysis period provides measurable emissions reductions and sufficient refurbishments for technology demonstration and maturation. This lowers the threshold for viable OOS business models, enabling technology demonstrations with measurable returns. As the technologies become increasingly mature, future constellations could be intentionally designed to be compatible with the CAAS CONOPs. Despite intuition suggesting distributed infrastructure would be necessary, a single warehouse captured much of the operational efficiency and environmental benefits of multi-depot architectures while avoiding multiplicative capital costs. Like a rain barrel collecting enough water in a rainstorm to water the garden, one warehouse in a large constellation collects enough satellites to provide sustainability benefits while moving the needle on space economy circularity.

With the assistance of cost-neutral policy mechanisms, the framework is able to increase refurbishment throughput and "catch more rain" without significant financial strains on either the government or the constellation operator. The analysis identified specific policy parameter combinations that achieve cost neutrality over the 30-year horizons while significantly increasing refurbishment throughput. Cost-neutral hybrid fee-rebate designs demonstrate that catalyzing sustainable space infrastructure is possible without imposing permanent fiscal burden on the government or the in-

dustry.

5.2 Critical Insights for Decision-Makers

This research reveals several critical insights for advancing on-orbit servicing and circular space economies:

First, policy should be developed in coordination with infrastructure planning. The framework demonstrates that policy interventions are effective when aligned with specific CONOPs and strategies. This interdependence suggests that future space sustainability regulations should emerge through collaborative processes involving both regulators and infrastructure developers rather than treating policy as CONOPs-agnostic.

The second critical insight is that reusable second-stage rockets emerge as transformative technology. The dual-mission architecture that leverages these vehicles for warehouse resupply and satellite return—never sending a rocket home empty-handed, avoiding "deadheading"—provides substantial NOx emission reductions while improving cost efficiency. The pathway to circular space economies need not rely solely on in-space manufacturing; strategic integration of Earth-to-orbit transportation can advance sustainability objectives while remaining economically competitive. Supporting circular space infrastructures and ADR requires high launch cadence, and with traditional second stages burning up in atmosphere, this risks contributing more atmospheric damage than it saves. Therefore, reusable second-stage vehicles are critical for environmental goals.

Meeting immediate business needs while considering long-term sustainability is key to enticing private sector participation. The first step toward circular economies lies in identifying services that are financially attractive to spacecraft operators in the near-term that offer dual-use potential for sustainability initiatives. ADR capabilities hosted in space-based infrastructures provide immediate economic value through

debris removal services and spare satellite hosting while simultaneously enabling collection operations that are essential for refurbishment. Collection hubs help replenish satellite constellations with spares and active satellites, reducing the need to launch new satellites from Earth. Similarly, RPO-capable satellites improve operational flexibility for constellation management by facilitating rendezvous operations required for servicing. Incremental steps that serve immediate purposes while setting up future opportunities are essential for economic viability.

Another takeaway is that RPO capability represents an immediate priority. While significant emphasis has been placed on designing satellites for refueling, this thesis identifies urgent need for improving LEO mobility through RPO-capable platforms. Although slower than high-thrust vehicles, electric-propulsion satellites with RPO provide the foundational capability for collection and aggregation operations essential to the CAAS concept. If private sector operators remain hesitant, government customers, such as NASA and the Department of Defense, can catalyze adoption through contracts prioritizing RPO-capable satellite designs. This represents a low-barrier form of intervention that government agencies can implement immediately.

Understanding the impact of emissions at scale remains critical. At the scale of refurbishment encouraged by this framework, making up relatively small constellation portion, the biggest driver of emissions reduction is not reduced deorbits, but the benefits from reusable vehicles and mission co-opting. For reduced deorbits to substantially impact atmospheric emissions, servicing infrastructure would need to extend satellite lifetimes across entire fleets. This underscores another critical insight: support from Earth remains essential and will continue to be, and these launches carry steep environmental costs. The framework identifies a pragmatic first step towards a more self-sustaining space environment rather than claiming immediate independence from Earth-based logistics.

5.3 Actionable Recommendations

5.3.1 For Policy Makers

Near-term recommendations (5-year horizon)

The analysis reveals three distinct policy pathways depending on societal priorities:

For a cost-neutral priority, implementation of orbital use fees with subsidies for collection, returning subsidy funds to satellite operators over time, provides an effective mechanism. Policy 1 with moderate orbital use fees provides a cost-neutral increase in refurbishment throughput because it accounts for the operational mechanics of CAAS concept and incentivizes sustainable operations like increased satellite collections. Critically, to avoid slowing down industry growth over long timelines, governments could return subsidies in shorter intervals rather than after the full analysis period used in these experiments. Returning subsidies in 10-15 year intervals addresses cash flow concerns while still providing the mechanisms to increase refurbishing throughput.

For a sustainability priority, subsidizing initial infrastructure deployment and implementing cost-neutral orbital use fee structure offers the strongest pathway. Policy 1 variations with initial warehouse subsidies and ongoing fees with rebates for collection provides the emissions savings of larger infrastructures while the cost-neutral fee structure incentivizes ongoing throughput. Upfront subsidies overcome coordination failures that prevent private investment in unproven infrastructure, while ongoing fee-rebate structure sustains operations without permanent fiscal burden.

For a balanced cost/sustainability priority, subsidizing one initial warehouse plus additional warehouses over time, contingent on demonstrated performance milestones, presents the optimal approach. This staged approach deploys initial infrastructure with subsequent warehouse subsidies triggered when refurbishment throughput, cost performance, or technology maturation reaches specified thresholds. Growth-focused

policies with flexible warehouse deployment over time proves more cost-effective than deploying all warehouses at once. While this does not provide the same immediate emissions savings as deploying multiple warehouses outright, it leverages the time value of money to reduce the experienced cost of large-scale infrastructure deployment while achieving meaningful sustainability gains.

All pathways should be paired with immediate satellite RPO capability and flexible satellite serviceability while the focus should prioritize standardized docking plates, autonomous navigation capabilities, and Earth-based refurbishment.

Long-term recommendations (10+ year horizon)

While not included within the scope of this thesis, policy makers could adopt adaptive policy frameworks that can pivot based on observed outcomes. A dynamic policy approach might start with a growth-focused plan emphasizing infrastructure deployment subsidies and minimal fees before transitioning to a cost-neutral policy after servicing infrastructure grows to a sufficient scale. The transition could be triggered by objective metrics like aggregate refurbishment throughput thresholds or the size of the subsidy fund.

Implementation considerations by jurisdiction

Direct implementation of orbital use fees in the United States faces political challenges, since the current climate emphasizes deregulation and commercial space competitiveness. However, opportunities exist for indirect pathways. The U.S. Space Force values dynamic responsiveness and maneuverability without regret, which are characteristics that are aligned with RPO-capable satellite architectures. Federal procurement contracts requiring RPO capability for defense satellites could mature the technology, creating spillover benefits for commercial applications. Additionally, orbital carriers being developed by companies like Gravitics for defense applications

share architectural similarities with the collection warehouses proposed in this thesis. While primarily intended to house defense assets, these platforms could be adapted with exterior docking plates to demonstrate collection capabilities in LEO, effectively piloting the CAAS concept under defense auspices before commercial scaling.

The U.S. could also leverage existing regulatory frameworks. The FCC's licensing authority over commercial satellites provides a mechanism for RPO capability requirements without new legislation. License conditions could mandate standardized docking interfaces and autonomous navigation capabilities as prerequisite for spectrum access, creating policy pressure through regulatory rather than fiscal means. Proceeds from FCC fines could help subsidize the relevant R&D initiatives.

The United Kingdom and European Space Agency are more likely than the U.S. to implement orbital use fees and establish the prescribed public-private partnerships. European precedent for environmental regulation, such as carbon pricing, Extended Producer Responsibility schemes, and Pigouvian taxation, provides the political and institutional foundation for similar space policies. Existing collaboration between Astroscale, OneWeb, and UK government to remove a failed OneWeb satellite demonstrates both technical capability and political willingness for public-private partnerships related to space sustainability. European implementation should emphasize cost-neutral designs to minimize competitiveness concerns.

Japan presents a favorable jurisdiction for comprehensive CAAS policy implementation. The Japanese space industry has demonstrated not only manufacturing capability in space sustainability technologies, but sustained governmental interest and financial commitment. Japan's coordinated industrial policy model, where government works closely with industry to develop strategic sectors, aligns naturally with the hybrid fee-rebate structures identified as most effective. Cultural and political emphasis on long-term planning and environmental stewardship, combined with technical leadership in robotics and precision manufacturing applicable to OOS systems,

positions Japan as a potential early leader in CAAS deployment.

Challenges of policy implementation

Several cross-cutting implementation challenges require attention regardless of jurisdiction. First, while cost-neutral policies achieve revenue neutrality over 30 year horizons, interim periods may see cash flow constraints for operators, particularly early entrants that bear fees before substantial subsidies become available. This cash flow mismatch could slow constellation deployment or deter new entrants. Incremental payouts from subsidy funds or exemptions for early entrants could mitigate this concern while maintaining long-term revenue-neutral structure.

Second, designing policy in coordination with technology and CONOPS development risks the appearance of "choosing winners" by favoring specific companies or technologies. Governments must maintain technology and vendor agnosticism to preserve competitive markets. However, the proposed policy approach is relatively agnostic regarding specific manufacturers: orbital use fees with rebates for any satellite collected to any certified facility, regardless of who operates the facility or builds the satellites. This opens the field for any entity to develop RPO-capable satellites or warehouse technology while providing a clear demand signal for making such investments.

5.3.2 For Satellite Constellation Operators

Near-term recommendations (5-year horizon)

Constellation operators should prioritize the development of electric propulsion satellites with RPO capability. This involves designing satellites with standardized docking plates, autonomous navigation sensors, and sufficient propellant reserves for terminal approach maneuvers. RPO-capable architectures could provide operational value beyond the applications considered in this thesis, such as improved collision avoid-

ance and the potential for satellite repositioning to optimize coverage during demand fluctuations.

Additionally, operators should begin demonstrating routine docking operations through low-risk pathways. Initial demonstrations could use the International Space Station or other emerging commercial space stations as docking targets and test hardware and software systems before deploying dedicated CAAS infrastructure. Rideshare missions could carry dedicated docking demonstration payloads, such as small experimental spacecraft that practice approach and berthing maneuvers with host vehicles or deployed docking targets.

Operators should also consider secondary revenue sources for collection capabilities beyond primary refurbishment mission. Collection hubs in valuable orbits could serve as temporary housing for defense or scientific payloads that require flexible deployment, effectively functioning as orbital "coworking" spaces. They could provide propellant resupply for deep-space mission vehicles using LEO as staging point, or serve as data relay nodes. These auxiliary services could improve the business case for deploying initial warehouse infrastructure before refurbishment demand fully materializes.

Finally, operators should deploy spare warehouses with the capability to dock with and refuel Active Debris Removal vehicles. This creates a symbiotic relationship: ADR vehicles gain operational flexibility from orbital refueling, reducing mission costs while warehouses gain immediate utility by serving ADR operations while collection infrastructure matures. This thesis finds that on-orbit refueling is more valuable for ADR vehicles than it is for collected satellites. This bootstrapping strategy provides near-term revenue while building toward longer-term refurbishment operations.

Long-term recommendations (10+ year horizon)

Operators should continuously evaluate environmental and industry conditions to determine the optimal mix of space-based versus Earth-based refurbishment. If Earth-return vehicle technology matures rapidly and costs decline as projected, Earth-based refurbishment becomes the strongly preferred pathway. In this case, operators should focus satellite design on ease of Earth-return and refurbishment, focusing on modular construction, accessible component bays, non-destructive disassembly features. Environmental testing should include return loads to ensure components survive round-trip transport.

Conversely, if Earth-return technology development encounters delays or cost floors prevent economic viability, space-based refurbishment may become relatively more attractive. In this case, operators should pivot toward making satellites fully serviceable in orbit by considering swappable components, robotic manipulation access, and diagnostic interfaces. The CAAS concept's value lies precisely in maintaining multiple options rather than over-committing to one refurbishment pathway too early.

5.3.3 For Satellite Manufacturers

Near-term recommendations (5-year horizon)

Manufacturers should explore incentive structures for the return of their own satellites. Companies could offer rebates or trade-in credits for constellation operators returning satellites at end-of-life. This could provide several benefits, such as (1) the opportunity to study long-term space environment effects on materials and components, (2) recovery of valuable materials, particularly rare earth elements in solar cells and electronics, (3) potential component reuse after testing and refurbishment, and (4) public relations value in demonstrating environmental stewardship. There are

terrestrial precedents for manufacturer take-back programs, such as automotive manufacturers offering trade-in values for end-of-life vehicles, Apple’s recycling programs for iPhones with credit toward new devices, and industrial equipment manufacturers refurbishing and reselling used machinery.

Long-term recommendations (10+ year horizon)

Over the long term, manufacturers should develop modular satellite architectures that are optimized for whichever refurbishment pathway emerges as dominant. If Earth-return proves viable, they should design for terrestrial disassembly, testing, and recertification. If space-based refurbishment prevails, they should design for robotic manipulation and component replacement in microgravity.

Manufacturers should also investigate strategic positioning in the OOS value chain. Vertically integrated manufacturers could leverage in-house refurbishment capabilities, capturing value from both initial production and subsequent lifecycle services. Alternatively, manufacturers might partner with specialized OOS providers, focusing on new production while licensees handle refurbishment. The optimal strategy depends on core competencies, capital availability, and the evolution of the market structure.

5.3.4 For On-Orbit Servicing Providers

Near-term recommendations (5-year horizon)

New entrants or pivoting companies should focus initial deployments on collection and aggregation capabilities and ADR vehicle refueling rather than sophisticated satellite servicing. Results from this thesis indicates that warehouses that provide berthing and temporary storage for satellites while providing dock and refueling services for ADR vehicles capture most of the environmental and economic value of full-service facilities. This suggests staged market entry is the best approach: deploy collection

capability first and validate operations and demand before adding satellite servicing capabilities.

Additionally, service providers should prioritize partnerships with reusable launch vehicle developers. The symbiotic relationship between launch vehicle operators getting secondary down-mass revenue for their fully reusable launch vehicles and OOS providers needing transport for collected satellites creates the potential for mutually beneficial partnerships. Joint ventures, long-term contracts, or vertical integration between launch and servicing could improve the economics for both parties.

Providers should also develop diverse revenue models. Pure refurbishment services may prove insufficient for financial viability, particularly during the early years. Auxiliary services, such as orbital storage for government or commercial payloads, data relay services, or hosting scientific payloads could provide a base revenue as the collection infrastructure scales up.

Long-term recommendations (10+ year horizon)

As the industry matures, service providers should develop sophisticated pricing strategies that capture the full value of the services provided. Initially, pricing may need to be penetration-level to attract early customers and prove concepts. However, as servicing becomes established and potentially required by regulation, providers can implement value-based pricing that reflects the full economic and regulatory value of debris removal, emissions reduction, and satellite life extension.

Service providers should also participate actively in standards development for OOS interfaces, protocols, and practices. First movers that establish de facto standards gain a lasting competitive advantage as satellite operators invest in compatibility with deployed infrastructure. However, overly proprietary approaches risk fragmentation and inefficiency. A balanced strategy of open standards for core interfaces with proprietary innovations in service delivery may optimize both adoption

and competitive positioning.

5.4 Lessons Learned and Recommendations for Future Researchers

Several methodological lessons emerged that will inform future work in this domain:

The use of high-performance computing clusters for large parametric studies is essential for recommended future work. This thesis evaluated dozens of distinct configurations across 80 Monte Carlo samples per test. More comprehensive parametric studies, particularly sensitivity analyses that systematically vary all input parameters across wider ranges, would provide a better idea of how assumed values and sources of uncertainty impact viability but would require significantly more runs. Future research should plan for cluster computing from the outset, designing simulation code for efficient parallelization and result aggregation.

Another lesson is that the numerical results of this thesis are use-case-specific while sensitivities and trends provide lasting insights. The specific numerical findings apply to the constellation parameters used in this analysis. Future researchers should emphasize identifying structural insights, threshold effects, and consistent orderings across scenarios rather than focusing on point estimates from particular parameter sets.

Overall, multi-domain uncertainty integration is essential but proved to be challenging. This framework integrated technological uncertainty (like Earth-return vehicle maturation), economic uncertainty (like launch costs, satellite manufacturing costs, and infrastructure costs), and operational uncertainty (like failure rates and technology obsolescence). Capturing all simultaneously stressed both modeling and computational resources, but was essential for meaningful conclusions. Isolating uncertainties would produce misleading results because interactions between configuration variables and sources of uncertainty were prominent.

Lastly, decision rule specification matters enormously. This thesis uses conditional-

go rule types with three-level decision hierarchy: individual satellite, satellite constellation fleet-level, and servicing infrastructure-level. The framework makes decisions at particular instances by considering the net present value of satellites, the state of uncertainty, and the current available options. These decision rules provide intuitive, straightforward logic modeling realistic human decision-making, though in some cases, decision multipliers were implemented and tuned through trial-and-error to perfect outcomes. While backward induction for real options analysis and multi-stage stochastic programming are powerful approaches, they're computationally intractable for this problem's complexity. Future work could explore more sophisticated optimization of these decision triggers and potentially consider validation against empirical data from analogous industries.

5.5 Future Work

This thesis establishes a foundation for analyzing CAAS and on-orbit servicing economics, yet several promising research directions could extend and refine this framework. These opportunities span technical modeling improvements, design considerations, market expansion, policy evolution, and operational architecture enhancements.

The first priority, which would not require major changes to the framework, should be to expand the computational budget using a cluster and run thousands rather than tens of scenarios. This would support development of accurate surrogate models and provide denser sampling of the design space, potentially revealing nonlinear interactions and edge cases not apparent in the current analysis. More data points could enable machine learning for decision rules, which could uncover non-intuitive decisions outperforming the human-tuned heuristics used in this thesis.

Incorporating optimization algorithms into the framework would improve both refurbishment throughput and operational efficiency. The current approach explores predefined configurations and policies while embedding optimization routines could

identify non-intuitive superior architectures emerging from the intersection of multiple decision variables. Similarly, optimizing decision rules for when to deploy flexible options, rather than using intuitive and manually-tuned decision triggers, could reveal more sophisticated deployment strategies that respond dynamically to system state and uncertainty realizations.

Another region for future work is improving cost modeling fidelity, which would strengthen confidence in the economic conclusions. While the cost models in this thesis capture major drivers including launch, manufacturing, and operations, higher-resolution modeling of refurbishment processes, logistics, and failure modes would enable more precise trade-space characterization. Collaboration with subject matter experts could refine these estimates beyond first-order approximations. Additionally, component-level failure modeling that differentiates between types of failures and requirements for servicing, rather than the simplified approach used here, would provide a more nuanced understanding of when different servicing capabilities provide value. This would inform capability prioritization decisions.

From a technical modeling perspective, refining the low-thrust trajectory model would better capture the benefits of collaborative maneuvering. The current implementation, which defaults to drift rate calculations for orbital transfers, likely underestimates efficiency gains achievable through coordinated low-thrust rendezvous operations. A higher-fidelity propagator that explicitly models thrust arcs, orbital mechanics, and propellant consumption would provide a more accurate assessment of collection and aggregation timelines, ultimately affecting both cost and throughput projections. Complementing these trajectory improvements, expanding environmental analysis beyond NOx emissions would provide a more comprehensive sustainability assessment. Full lifecycle analysis incorporating satellite manufacturing, launch vehicle production, ground operations, and end-of-life disposal would contextualize relative environmental benefits of refurbishment versus new satellite deployment. Under-

standing how returning satellites to Earth for refurbishment, rather than deorbiting them, affects total environmental footprint requires detailed lifecycle modeling that accounts for manufacturing energy, material extraction, and reprocessing capabilities.

Shifting focus to design considerations, longer operational lifespans deserve attention. Previous on-orbit servicing frameworks have examined this, but this thesis focused on servicing present-day LEO satellite designs. Understanding how constellation design might evolve if satellites are intentionally built for extended lifespans with servicing availability could unlock significant additional value. Building on this design perspective, the framework could guide constellation design for serviceability more directly. New entrants could distinguish themselves from existing mega-constellations by intentionally designing constellations with fewer, longer-lived, higher-value satellites positioned strategically to leverage CAAS from the start. Researchers could leverage this framework to inform these strategic design decisions.

Regarding market expansion, opening the framework to multi-constellation servicing with different satellite sizes and types could dramatically expand the business case, despite coordination challenges involved. This thesis focused on a single owner-operator constellation, but realistic OOS infrastructures would likely serve multiple customers. The economic benefits of shared infrastructure across multiple operators could fundamentally alter the value proposition.

Policy considerations present particularly rich opportunities for future investigation, approached from two complementary perspectives. From the policymaker perspective, exploring time-varying policy implementation merits investigation. While this thesis evaluated static policy schemes, future work could examine deliberate policy transitions over mission timelines. Some policies excel at enabling infrastructure growth while others optimize cost-neutral throughput once infrastructure exists. For instance, starting with growth-focused subsidies before transitioning to cost-neutral fee-rebate structures after certain infrastructure milestones could unlock significant

additional value while managing fiscal burden over time. From the operator perspective, treating policy as time-varying uncertainty presents a complementary research direction. This would examine how constellation operators should account for policy uncertainty when they don't control what policies get implemented or when they change. If policy represents exogenous uncertainty rather than known regulation, how should operators design flexible business strategies and infrastructure investments that remain viable across different policy scenarios? This would inform resilient constellation design and operational strategies.

The bidirectional relationship between CAAS and reusable second stages warrants deeper study. CAAS relies on returning satellites to Earth, but simultaneously creates a business case for reusable vehicles by making return trips valuable. Understanding this mutually reinforcing dynamic deserves dedicated investigation, particularly as launch vehicle reusability continues maturing.

Finally, extending the CAAS concept of operations could unlock operational efficiencies not explored in the current framework. Two promising directions merit investigation. First, ADR vehicles employing electric propulsion rather than chemical propulsion deserve consideration. Because the warehouse/CAAS system inherently operates on slower timescales by using J2 drift for orbital plane changes rather than impulsive maneuvers, there is less premium on rapid orbital transfers. Electric propulsion systems, while slower, offer substantially higher specific impulse and lower propellant mass requirements. ADR vehicles could spiral between collection targets rather than executing fast Hohmann transfers, potentially reducing operational costs and extending vehicle operational lifetimes. The trade between transfer time and propellant efficiency becomes more favorable when collection operations already accommodate drift-based timescales spanning weeks to months.

Second, reusable second stages serving dual roles as orbital transfer vehicles or mobile servicers could fundamentally alter the operational architecture. Rather than

viewing reusable stages solely as Earth-return transportation, they could function as agile orbital assets capable of traveling through LEO independently. This would enable faster satellite collection, distributed warehouse locations, or even servicing capabilities without requiring separate ADR vehicle fleets. The bidirectional nature of reusable stages makes them natural candidates for orbital maneuvering roles that capitalize on their propulsion systems during the on-orbit portion of missions. Understanding how to optimize this dual-use capability could reveal synergies between launch vehicle reusability and on-orbit servicing infrastructure that strengthen the business case for both.

5.6 Closing Perspective

This thesis demonstrates that the pathway to circular space economies emerges not from technological breakthroughs alone, but from systematic understanding of how technical architectures, flexible deployment strategies, and policy mechanisms interact across uncertain futures. Circular space economies emerge from systematic understanding of not just technology, but interactions between engineering, business strategy, and policy—designing for one in a vacuum won’t achieve circularity.

The framework reveals consistent patterns indicating viable pathways exist depending on priorities. RPO capability, reusable launch vehicles, warehouse-deployed ADR, and strategic policy-flexibility pairing consistently improve outcomes relative to baseline approaches. There are trade-offs, but the framework provides decision-makers with a screening tool to identify strategies balancing economic viability with sustainability objectives.

The key insight is that viability comes through deliberate infrastructure development that responds to immediate needs while providing opportunities for sustainable operations in the future. For private sector appeal, de-risking these steps as much as possible is critical. This thesis provides pragmatic vision that shows achievable

first steps. The transition to sustainable space operations need not follow an all-or-nothing approach; starting with just the portion of the constellation that reaches the breakeven point offers a pragmatic pathway for encouraging private companies to adapt larger operations. If 30-year deployment results in 4-7 operational warehouses, subsequent decades will benefit from established infrastructure capable of higher refurbishment rates.

As humanity's presence in LEO intensifies with projections of tens of thousands of satellites in coming decades, the need for a circular space economy becomes increasingly urgent. This research demonstrates technical pathways, strategic approaches, and policy frameworks that aim to make sustainable LEO operations achievable. The question is no longer whether circular space economies are possible, but whether stakeholders possess the collective will to implement them before unsustainable practices become entrenched and environmental challenges blossom.

The next generation of space operations becomes achievable when we align technical capabilities, business strategy, and policy perspectives. These findings suggest that viable circular space economies in LEO are achievable through deliberate, incremental infrastructure development supported by aligned policy interventions - a more pragmatic vision than revolutionary transformation, but one grounded in the economic and technical realities that will shape the next generation of space operations.

Appendices

APPENDIX A

ASSUMPTIONS, EQUATIONS, AND MODELS

Table A.1: Use Case Simulation Parameters

Variable Name	Range
number of planes	18
satellites per plane	36
customer altitude	1200km
parking orbit altitude	796km
inclination	86.4 degrees
number of in-plane spares	2 satellites
satellite dry mass	150 kg [192]
satellite fuel mass	12 kg
dry mass ADR	150 kg [80]
fuel mass ADR	150 kg
ISP ADR	230 s
warehouse max capacity	35 satellites
warehouse initial capacity	5 spare satellites
warehouse electric fuel mass	500 kg
warehouse chemical fuel mass	1000 kg
warehouse dry mass	1000 kg
Xenon cost	\$5000/kg [193]
Green monopropellant cost	\$100/kg [194] [195]
discount rate	0.03

Table A.2: Rocket Performance Parameters

Parameter	Falcon 9	Starship	Electron	Neutron	Stoke Nova	Terran R	New Glenn
ISP_1	311[227]	330	311[228]	330	345[229]	330	340
mass_stage_full [kg]	549 054	5 000 000	13 000	480 000	100 000	1 500 000	1 600 000
ISP_2	348[227]	380	343[228]	360	425[229]	375	445
payload_max_capacity [kg]	22 800	100 000	300	13 000	3000[229]	23 500	45 000
V_stage_sep [m/s]	3410[230]	2300	2333[231]	2500	2500	2500	8000
second_stage_mass [kg] Equation 3.34	180 000	2 460 000	6050	222 000	47 800	693 000	226 796
V_final [m/s]	7800	7800	7800	7800	7800	7800	7800
Mdry2 [kg] Equation 3.35	26 800	460 000	890	36 400	10 400	140 500	27 215.5
initial launch cost (\$/kg)	3986[200]	100	23 437[236]	3846[237]	250	2340[238]	1511 [200]
estimated years until mature	0	5	0	4	5	5	3

322

Table A.3: Sources of Uncertainty and Modeling Parameters

Uncertain Quantity	Model/Method	Initial Value(s)	Uncertainty Parameters	Eq.
Constellation Revenue	Geometric Random Walk / Log-normal PDF	\$1.4B/year [196]	$\alpha = 0.059$ [197], $\sigma = 0.15$	Equation A.1
Launch Cost	Log-normal PDF with volatility cone	α_m based off present launch cost per kg and Citi Bear Case [198]	$\sigma = 0.1$ before 2040, $\sigma = 0.35$ and $\alpha_m = 0$ after 2040	Equation A.1

Continued on next page

Table A.3 – continued from previous page

Uncertain Quantity	Model/Method	Initial Value(s)	Uncertainty Parameters	Eq.
Launch Delay	Processing + Exponential	$T_{\text{processing}} = 3$ months, $\mu_{\text{launch}} = 2$ months	Exponential with μ_{launch}	Equation A.2
ADR Launch Delay	Learning Curve + Exponential	$\text{min}_{\text{time}} = 0.5$ years, $\text{initial}_{\text{extraTime}} = 1.5$ years, $\lambda = 0.2303$	Exponential with μ_{launch}	Equation A.3
Satellite Manufacturing Cost	Log-normal PDF	\$900,000 [191]	$\alpha_m = -0.075$ [199], $\sigma_m = 0.1$	Equation A.1
ADR Vehicle Cost	Learning Curve + Time-based decay	\$48M (initial) [101]	$\lambda \sim U(0.1, 0.5)$, $P_{\text{min}} \sim U(0.5, 0.8) \times \text{cost}_i$, $r = 0.03$	Equation A.4
ADR Collection Success	Learning Curve (Power Law)	$S_0 \sim U(0.7, 0.99)$	$\lambda \sim U(0.05, 0.25)$	Equation A.6
Space-Based Services (ADR Operation, Refuel, Repair, Obsolete Repair)	Log-normal PDF	\$250k / \$250k / \$562.5k / \$687.5k	$\alpha = -0.0375$, $\sigma = 0.1$	Equation A.1

Continued on next page

Table A.3 – continued from previous page

Uncertain Quantity	Model/Method	Initial Value(s)	Uncertainty Parameters	Eq.
Warehouse Cost	Learning Curve	\$100M (initial)	$\lambda \sim U(0.1, 0.5)$, $P_{\min} \sim U(0.5, 0.8) \times \text{cost}_i$, $N_{\max} = 54$	Equation A.7
Warehouse Maintenance Cost	Lognormal PDF (cost)	Every 15 years: \$5M median repair cost	Cost: $\mu = \ln(5 \times 10^6)$, $\sigma = 1.0$	Equation A.8
Earth-Based Services (Repair, Refurbish, Obsolete Repair)	Log-normal PDF	\$450k / \$200k / \$550k (repairable); \$855k (non-repairable)	$\alpha = -0.075$, $\sigma = 0.1$	Equation A.1
Warehouse Upgrades	Log-normal PDF	\$8M / \$15M	$\alpha = -0.0375$, $\sigma = 0.1$	Equation A.1
Satellite Upgrade R&D	Log-normal PDF	\$5M / \$9M R&D + 6% / 10% cost increase	$\alpha = -0.0375$, $\sigma = 0.1$ on R&D	Equation A.1
Technology Obsolescence	Weibull Utility Function	N/A	$k = 2$, $\lambda = 1$, $\beta = 2$	Equation A.9
Failure Time	Exponential (from MTBF) and Failure Change Rate, Uniform	MTBF starting with 1 failure/year, includes exponential change rate with random parameter from uniform distribution	MTBF Derived from fleet size, exponential failure change rate varies uniformly between 0.005 and 0.08	Equation A.10, Equation A.11, and Equation A.12

Continued on next page

Table A.3 – continued from previous page

Uncertain Quantity	Model/Method	Initial Value(s)	Uncertainty Parameters	Eq.
Failure Type	Bernoulli Trial	50% inoperable	None (uniform draw)	—
Collision Event Cost	Exponential + Uniform + Cascade Probability	\$10k base	$\beta = 0.01$, $P_{\text{cascade}} = \min(0.02T, 0.5)$	Equation A.13

Table A.4: Initial Cost Assumptions (Year: 2025)

Variable Name	Value	Justification
Satellite Unit Cost	\$900,000	Based on estimated OneWeb satellite cost [191]
Earth Refurbishment (General)	\$200,000	Includes labor, testing, parts and upgrades, and facilities. It's assumed that refurbishing is roughly 20-40% of manufacturing a new satellite, based on the economics of refurbishing reusable rockets, which is 65% cheaper than launching new [221]
In-space/Earth Repair (Repairable Satellite Failure)	\$450,000	More effort than refurbishment, roughly half the cost of a new satellite
In-space/Earth Repair (Obsolescence, Repairable)	\$550,000	More expensive than repair due to payload upgrade and related testing

Table A.4 – continued from previous page

Variable Name	Value	Justification
Earth Repair (Obsolescence, Non-repairable)	\$855,000	Requires more extensive labor to swap a payload in a satellite not meant to be repaired, roughly the same cost as a new satellite
Satellite Refuel Upgrade, R&D	\$5,000,000	Redesign for mechanical and fluid interface, propulsion system adjustments, electrical & software integration, testing and certification
Satellite Repair Upgrade, R&D	\$9,000,000	Design modularization, standardized interfaces, on-board failure monitoring, testing, qualification
Warehouse Refuel Upgrade	\$8,000,000	Xenon storage and pressure management, refueling system with specialized docking, precision control systems, diagnostic monitoring, coordination between depot and satellite, training, certification, and testing
Warehouse Repair Upgrade	\$15,000,000	Diagnostic systems and related sensors, robotic capability and tooling, spare parts, power systems for satellite battery test & recharge, specialized software, training, and testing
ADR Vehicle Cost	\$48,000,000	Based on \$16 million (USD) Astroscale contract that covers about 1/3 of vehicle cost [101]

Table A.4 – continued from previous page

Variable Name	Value	Justification
ADR Operation	\$250,000	Includes the cost for ground support, GNC, RPOD, system checks, mission operations
ADR Operation for RPO satellite	\$83,250	Includes the cost for ground support, GNC, RPOD, system checks, mission operations. Less than classic ADR operations because RPO satellites are upgraded to be cooperative and maneuver duration is shorter
In-Space Refueling Operation	\$250,000	Includes the cost for ground support, controls, robotic operation, system checks, depreciation of hardware, labor, mission operations
Warehouse Cost	\$100,000,000	Based on analogous servicing spacecraft budgets, such as DARPA RSGS [222] and MEV-2 [223], which have more capability than initial spare warehouse configuration. Can also be compared to the scale of large GEO satellites, such as Intelsat 10-02 [223], factoring out launch cost to GEO
Deorbit Cost (From constellation altitude)	initially \$100,000; set to 1/9 of space-craft cost and varies proportionally	Based on NASA estimate of additional cost for medium satellite's extra propellant to immediately deorbit from 800 km ranges (between \$85,000 and \$425,000 [224])

Continued on next page

Table A.4 – continued from previous page

Variable Name	Value	Justification
Satellite to Warehouse RPO Cost	1/2 of present deorbit cost	conservative estimate based on estimated deorbit cost [224])
Deorbit Cost (From parking orbit altitude)	1/2 of present deorbit cost	conservative estimate based on estimated deorbit cost [224])

Table A.5: Latin Hypercube DOE Parameters for Uncertain Variables

Parameter	Min	Max
Earth-based repair operation cost streamlining cost multiplier	0.5	1
Earth-based refuel operation streamlining cost multiplier	0.5	1
Technology obsolescence intensity parameter	1	5
Time to technology obsolescence onset (years)	5	15
Active Debris Removal (ADR) initial success rate	0.7	0.99
ADR success rate learning curve exponent	0.05	0.25
Warehouse cost learning curve exponent	0.1	0.5
Warehouse cost minimum cost fraction	0.5	0.8
ADR cost learning curve exponent	0.1	0.5
ADR minimum cost fraction	0.5	0.8
Satellite cost multiplier for RPO-capability upgrade	1	1.5
Satellite cost multiplier for repair upgrade	1	1.3
Satellite cost multiplier for refuel upgrade	1	1.2
Satellite return to Earth cost (multiplier of present launch cost)	0.2	0.99
Simulation year that returning satellites to Earth becomes possible	7	15
Satellite failure rate change rate (per year)	0.005	0.079
Percentage of launch vehicles with reusable 2nd stages	20%	70%
Reusable 2nd stage fleet percent multiplier if Earth-return available	1×	1.25×

A.0.1 Uncertainty Equations

Log-normal probability density function for the geometric random walk method

The geometric random walk method is commonly used to model revenue and was used in previous OOS flexibility frameworks [10, 141, 142].

$$p_{\tau}^{(m)}(x) = \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma_m \sqrt{\tau}} \frac{1}{x} \exp \left\{ -\frac{(\ln(x) - (\alpha_m - \sigma_m^2/2)\tau)^2}{2\sigma_m^2 \tau} \right\} \quad (\text{A.1})$$

In the equation above, α_m is drift, ω_m is volatility, and τ is time.

Launch Cost, Delay, and Payload

Launch cost scenarios use log-normal models with volatility cones based on Citi Research projections [198] with a sigma m approximated at 0.1. After year 2040, sigma m increases to 0.35 since this goes beyond the range of the citi research projections.

Uncertainty is also applied to launch time, with $T_{\text{processing}} = 3$ months and $\mu_{\text{launch}} = 2$ months such that

$$T_{\text{launch}} = T_{\text{processing}} + X_{\text{launch}}, \quad X_{\text{launch}} \sim \text{Exponential}(\mu_{\text{launch}}) \quad (\text{A.2})$$

For ADR mission launched directly from Earth, $T_{\text{processing}}$ follows a learning curve as more ADR vehicles are manufactured and launched.

$$\text{adr}T_{\text{processing}} = \min_{\text{time}} + (\text{initial}_{\text{extraTime}} \times e^{-\lambda t}) \quad (\text{A.3})$$

where $\min_{\text{time}} = 0.5$ years, $\text{initial}_{\text{extraTime}} = 1.5$ years, and $\lambda = 0.2303$.

This framework assumes that downmass payload capacity on reusable second stages is 1/3 of its launch payload capacity.

ADR Vehicle Cost

ADR vehicle costs are modeled using a time-based exponential decay starting from an initial cost of \$48,000,000 [101]. The cost incorporates a learning curve effect based on the number of vehicles manufactured:

$$n = \{1, 2, \dots, N_{\text{max}}\}; \text{value}(n) = P_{\text{min}} + (1 - P_{\text{min}}) \cdot e^{-\lambda n} \quad (\text{A.4})$$

where n represents the number of ADR vehicles added (starting from 1), N_{\max} is the maximum number of ADR vehicles considered (assumed to be equal to the number of planes, 18), P_{\min} is the minimum ADR cost fraction (asymptotic limit), which is randomly sampled between 0.5 and 0.8, and λ is the learning exponent controlling the rate of decay, which is randomly sampled between 0.1 and 0.5.

Additionally, a time-based cost reduction applies with annual reduction rate $r = 0.03$ and time t in years:

$$\text{cost_reduction} = (1 - r)^t \quad (\text{A.5})$$

Space-Based Operations and Costs

ADR vehicles are given a chance of successfully collecting a satellite based on a learning curve trend which is given as follows:

$$n = \{1, 2, \dots, N_{\max}\}; \text{value}(n) = 1 - (1 - S_0) \cdot n^{-\lambda} \quad (\text{A.6})$$

In this equation, n is the number of satellites collected (starting from 1), N_{\max} is the maximum number of collected satellites considered (assumed to be 50), after which point the collection probability remains the same, S_0 is the starting point (initial value at $n = 1$), which is a random value that ranges between 0.7 and 0.99, and λ is the learning exponent (controls the rate of improvement), which is a random value that ranges between 0.05 and 0.25

The cost of each ADR operation is assumed to start at \$250,000 and this value changes over time following the same geometric random path as the other space-based operations. These space-based services are proportional to the same random path because it is assumed that they are all related to the same general capabilities. The ADR operation cost for RPO satellites is one third the typical ADR operation cost (\$83,250 per operation) because RPO satellites are upgraded to be cooperative

targets, and the maneuver duration is shorter due to this enhanced capability.

Warehouse and Depot Cost Uncertainty

Warehouse cost uncertainty mirrors the ADR vehicle cost modeling approach, with an initial cost of \$100,000,000 that decreases with each new warehouse added to the system. The learning-curve-based reduction follows:

$$n = \{1, 2, \dots, N_{\max}\}; \text{value}(n) = P_{\min} + (1 - P_{\min}) \cdot e^{-\lambda n} \quad (\text{A.7})$$

where n is the number of warehouses added (starting from 1), N_{\max} is the maximum number of warehouses considered, P_{\min} is the minimum performance level (asymptotic limit), and λ is the learning exponent (rate of decay) randomly sampled between 0.1 and 0.5.

Warehouse Failure Cost Modeling In addition to initial construction and operational costs, warehouses require maintenance every 15 years. The cost associated with each warehouse repair is modeled using a lognormal distribution to capture the high variability and right-skewed nature of major infrastructure repair costs:

$$C_{\text{failure}} \sim \text{Lognormal}(\mu = \ln(5 \times 10^6), \sigma = 1.0) \quad (\text{A.8})$$

where the median failure cost is approximately \$5 million, while $\sigma=1.0$ allows for significant variability. This modeling approach reflects the reality that warehouse infrastructure failures can range from minor system repairs to complete facility reconstruction, depending on the nature and severity of the failure event.

The total warehouse-related costs over the mission lifetime thus include initial construction costs (with learning curve and time-based reductions), operational costs, and stochastic failure repair costs that occur throughout the warehouse operational life.

Technology Obsolescence

This framework uses a Weibull-based utility function to reduce satellite revenue after it has reached the time of obsolescence, based the utility function $u(t)$, defined by on Geng et. al. [201]. The Geng model uses 3-parameter Weibull distribution with $\beta = 2$. The intensity metric is determined randomly for each scenario and time to obsolescence for each scenario comes from a Weibull distribution with shape, $k = 2$ and scale, $\lambda = 1$.

$$u_i(t) = u_{o,i} e^{-\left(\frac{(t-T_{obs,i})^\beta}{\theta_{obs,i}}\right)} \quad \text{for } t \geq T_{obs,i}, \quad (\text{A.9})$$

$$u_{\text{total}} = \sum_i u_i(t).$$

Failure Rates

OneWeb has experienced 4 satellite failures in 4 years [202] [203]. This paper therefore sets failure rate to 1 satellite fail per year, which is converted to mean time between failure (MTBF) using the following equation:

$$MTBF = ((1/(failureRate/(numPlanes * totalSatellitesPerPlane))) \quad (\text{A.10})$$

Every satellite is assigned a time to failure based on the MTBF:

$$T_{\text{failure}} \sim \text{Exponential}\left(\frac{1}{MTBF}\right) \quad (\text{A.11})$$

While every simulation starts with 1 failure per year, this failure rate decreases

over time, such that:

$$f = f_0 - \min(0.1 \cdot e^{r \cdot t}, 0.99) \quad (\text{A.12})$$

f = current failure rate

f_0 = original failure rate

Where:

r = failure rate change rate

t = current time

In this equation, failure rate change rate is a random value sampled from a uniform distribution, ranging from 0.005 to 0.08.

Cost of Collision Avoidance and Collision Events

$$P_{\text{collision}} = 1 - e^{-\beta \cdot N_{\text{fail}} \cdot T} \quad (\text{A.13})$$

In this equation, N_{fail} is the number of failed satellites, T is the number of years, and $\beta = 0.01$, the collision hazard rate per satellite per year. The base cost of a collision-related event is defined as $C_{\text{base}} = 10,000$. The probability of a cascade-type catastrophic event is:

$$P_{\text{cascade}} = \min(0.02 \cdot T_{\text{cascade}}, 0.5) \quad (\text{A.14})$$

The collision cost is then drawn from a uniform distribution, based on whether a cascade occurs:

$$C_{\text{collision}} = \begin{cases} C_{\text{base}} \cdot U(1000, 10000), & \text{with probability } P_{\text{cascade}} \\ C_{\text{base}} \cdot U(1, 1000), & \text{with probability } 1 - P_{\text{cascade}} \end{cases} \quad (\text{A.15})$$

The actual cost incurred is:

$$C = \begin{cases} C_{\text{collision}}, & \text{if a collision occurs (with probability } P_{\text{collision}}) \\ 0, & \text{otherwise} \end{cases} \quad (\text{A.16})$$

APPENDIX B

DECISION RULE ALGORITHMS

B.1 Satellite-by-Satellite Decisions




Figure B.1: Appendix: Satellite-Level Decision Tree Logic

B.2 Satellite-Constellation-Level Decisions

B.2.1 Upgrade satellites to be refuelable or repairable




Figure B.2: Appendix: Satellite Refuelability/Repairability Decision Tree

B.2.2 Upgrade satellites to be RPO-capable




Figure B.3: Appendix: RPO Upgrade Decision Rule Tree

B.3 Infrastructure-Level Decisions

This section presents the mathematical formulation of servicing infrastructure-level decision algorithms

B.3.1 Add a new warehouse




Figure B.4: Appendix: Add Warehouse Decision Rule Tree

B.3.2 Upgrade a warehouse




Figure B.5: Appendix: Upgrade Warehouse Decision Rule Tree

REFERENCES

- [1] Aerospace Corporation, *CORDS Reentry Database*.
- [2] FCC, “In the Matter of Space Innovation Mitigation of Orbital Debris in the New Space Age,” Tech. Rep., 2022.
- [3] M. A. Luu and D. E. Hastings, “On-Orbit Servicing System Architectures for Proliferated Low-Earth-Orbit Constellations,” *Journal of Spacecraft and Rockets*, vol. 59, no. 6, pp. 1946–1965, Aug. 2022.
- [4] K. Ikeya and K. Ho, *Analysis of Cooperative and Non-Cooperative Architectures for Multi-Plane On-Orbit Refueling*. Feb. 2023.
- [5] R. Buchs and T. Bernauer, “Market-based instruments to incentivize more sustainable practices in outer space,” *Current Opinion in Environmental Sustainability*, vol. 60, p. 101 247, 2023.
- [6] M.-A. Cardin, A. Mijic, and J. Whyte, “Flexibility and Real Options in Engineering Systems Design,” in *Handbook of Engineering Systems Design*, A. Maier, J. Oehmen, and P. E. Vermaas, Eds., Cham: Springer International Publishing, 2020, pp. 1–29, ISBN: 978-3-030-46054-9.
- [7] J. Saleh, G. Mark, and N. Jordan, “Flexibility: A Multi-disciplinary Literature Review and a Research Agenda for Designing Flexible Engineering Systems,” *Journal of Engineering Design - J ENGINEERING DESIGN*, vol. 20, pp. 307–323, Jun. 2009.
- [8] J. H. Saleh, G. Mark, and N. C. Jordan, “Flexibility: a multi-disciplinary literature review and a research agenda for designing flexible engineering systems,” *Journal of Engineering Design*, vol. 20, no. 3, pp. 307–323, Jun. 2009.
- [9] J. H. Saleh, E. S. Lamassoure, D. E. Hastings, and D. J. Newman, “Flexibility and the Value of On-Orbit Servicing: New Customer-Centric Perspective,” *Journal of Spacecraft and Rockets*, vol. 40, no. 2, pp. 279–291, Mar. 2003.
- [10] E. S. Lamassoure, “A Framework to Account for Flexibility in Modeling the Value of On-Orbit Servicing for Space Systems,” Ph.D. dissertation, MIT, Jun. 2001.
- [11] O. de Weck, R. de Neufville, and M. Chaize, “Staged Deployment of Communications Satellite Constellations in Low Earth Orbit,” *Journal of Aerospace Computing Information and Communication - J AEROSP COMPUT INF COMMUN*, vol. 1, pp. 119–136, Mar. 2004.

- [12] J. Lin, O. de Weck, R. de Neufville, and H. K. Yue, “Enhancing the value of offshore developments with flexible subsea tiebacks,” *Journal of Petroleum Science and Engineering*, vol. 102, pp. 73–83, 2013.
- [13] J. Rainbow, *FCC fines Dish Network for botched satellite de-orbit*, Oct. 2023.
- [14] P. C. M. Jakob, K. Ho, S. Shimizu, and S. Yoshikawa, “Optimal satellite constellation spare strategy using multi-echelon inventory control with stochastic demand and lead times,” *ArXiv*, vol. abs/1807.02685, 2018.
- [15] E. Levin, J. Pearson, and J. Carroll, “Wholesale debris removal from LEO,” *Acta Astronautica*, vol. 73, pp. 100–108, 2012.
- [16] R. Lucken, N. Hubert, and D. Giolito, “Systematic space debris collection using Cubesat constellation Share My Space, Orbit cleaning and Space Debris Recycling,” Jul. 2017.
- [17] M. Bowne, T. Sarton du Jonchay, and D. Mavris, “Collection-As-A-Service to Incentivize On-Orbit Servicing in LEO,” 2025.
- [18] M. Bowne, T. Sarton du Jonchay, and D. Mavris, “Analyzing the Impact of Policy within a Flexibility Framework for LEO-Based ISAM with Collection-As-A-Service, ASCEND Manuscript,” in *AIAA ASCEND 2025: Policy and Strategies for the New Space Economy*, Las Vegas: AIAA, Jul. 2025.
- [19] Eurocontrol, “Satellites to be Built and Launched,” Euroconsult, Courbevoie, Tech. Rep., Dec. 2022.
- [20] T. Sarton du Jonchay, “FRAMEWORK FOR THE DESIGN AND OPERATIONS OF SUSTAINABLE ON-ORBIT SERVICING INFRASTRUCTURES DEDICATED TO GEOSYNCHRONOUS SATELLITES,” Ph.D. dissertation, Georgia Institute of Technology, Atlanta, Dec. 2022.
- [21] C. Daehnick, J. Gang, and I. Rozenkopf, “Space launch: Are we heading for oversupply or a shortfall?” *McKinsey*, Apr. 2023.
- [22] R. Smith, “The Little-Known Secret That Could Cost Elon Musk \$8.2 Billion a Year,” *The Motley Fool*, 2024.
- [23] A. C. Boley and M. Byers, “Satellite mega-constellations create risks in Low Earth Orbit, the atmosphere and on Earth,” *Scientific Reports*, vol. 11, no. 1, p. 10642, 2021.
- [24] A. Lawrence *et al.*, “The case for space environmentalism,” *Nature Astronomy*, vol. 6, no. 4, pp. 428–435, 2022.

- [25] K. L. Jones and A. K. Jain, “The Green Circularity: Life Cycle Assessments in the Space Industry,” Aerospace Corporation, Segundo, Tech. Rep., Mar. 2023.
- [26] X.-S. Yap, J. Heiberg, and B. Truffer, “The emerging global socio-technical regime for tackling space debris: A discourse network analysis,” *Acta Astronautica*, 2023.
- [27] A. Nelson, “NATIONAL ORBITAL DEBRIS IMPLEMENTATION PLAN,” ORBITAL DEBRIS INTERAGENCY WORKING GROUP SUBCOMMITTEE ON SPACE WEATHER, SECURITY, and HAZARDS, Washington, D.C., Tech. Rep., Jul. 2022.
- [28] Space Safety Council, “Best Practices for the Sustainability of Space Operations,” Space Safety Council, Tech. Rep., Apr. 2023.
- [29] FCC, “In the Matter of Mitigation of Orbital Debris in the New Space Age,” Tech. Rep., 2020.
- [30] A. Jain and D. E. Hastings, “Global Climate Effect from Space Debris Reentry: Engineering and Policy Implications,” in *ASCEND 2023*, ser. ASCEND, American Institute of Aeronautics and Astronautics, Oct. 2023.
- [31] T. J. Colvin, J. Karcz, and G. Wusk, “Cost and Benefit Analysis of Orbital Debris Remediation,” NASA Office of Technology, Policy, and Strategy, Washington, D.C., Tech. Rep., Mar. 2023.
- [32] J. N. Pelton, *New Solutions for the Space Debris Problem*. Cham: Springer International Publishing, 2015, ISBN: 978-3-319-17150-0.
- [33] D. W. Keith, “GEOENGINEERING THE CLIMATE: History and Prospect,” *Annual Review of Energy and the Environment*, vol. 25, no. 1, pp. 245–284, Nov. 2000.
- [34] E. A. Parson and D. W. Keith, “End the Deadlock on Governance of Geoengineering Research,” *Science*, vol. 339, no. 6125, pp. 1278–1279, Mar. 2013.
- [35] M. K. Brown, H. G. Lewis, A. J. Kavanagh, and I. Cnossen, “Future Decreases in Thermospheric Neutral Density in Low Earth Orbit due to Carbon Dioxide Emissions,” *Journal of Geophysical Research: Atmospheres*, vol. 126, no. 8, e2021JD034589, Apr. 2021.
- [36] R. G. Ryan, E. A. Marais, C. J. Balhatchet, and S. D. Eastham, “Impact of Rocket Launch and Space Debris Air Pollutant Emissions on Stratospheric Ozone and Global Climate,” *Earth’s Future*, vol. 10, no. 6, e2021EF002612, Jun. 2022.

- [37] E. J. L. Larson, R. W. Portmann, K. H. Rosenlof, D. W. Fahey, J. S. Daniel, and M. N. Ross, “Global atmospheric response to emissions from a proposed reusable space launch system,” *Earth’s Future*, vol. 5, no. 1, pp. 37–48, Jan. 2017.
- [38] C. Park and J. V. Rakich, “Equivalent-cone calculation of nitric oxide production rate during space shuttle re-entry,” *Atmospheric Environment (1967)*, vol. 14, no. 8, pp. 971–972, 1980.
- [39] FAA, *Environmental*, Jun. 2023.
- [40] S. Fleming, *Circular economy examples - how IKEA, Burger King, Adidas and more are investing in a circular economy*, Dec. 2020.
- [41] NREL, *Pathways To Achieve New Circular Vision for Lithium-Ion Batteries*, Feb. 2021.
- [42] M. Gaskill, *NASA Achieves Water Recovery Milestone on International Space Station*, Jun. 2023.
- [43] European Space Policy Institute (ESPI), “On-orbit Servicing, Assembly, and Manufacturing: State of Play and Perspectives on Future Evolutions,” European Space Policy Institute (ESPI), Vienna, Tech. Rep., Oct. 2023.
- [44] R. B. Friend, “Orbital Express program summary and mission overview,” in *Proc.SPIE*, vol. 6958, Apr. 2008, p. 695 803.
- [45] M. Luu and D. E. Hastings, “Review of On-Orbit Servicing Considerations for Low-Earth Orbit Constellations,” in *ASCEND 2021*, ser. ASCEND, American Institute of Aeronautics and Astronautics, Nov. 2021.
- [46] S. Erwin, *Space Force eager to harness satellite-servicing technologies*, Jun. 2023.
- [47] Astroscale, *CisLunar Industries, Astroscale U.S. and Colorado State University Awarded \$1.7M Contract by the U.S. Space Force*, 2023.
- [48] A. H. Inc., *Astroscale’s ELSA-d Finalizes De-Orbit Operations Marking Successful Mission Conclusion*, Nov. 2024.
- [49] M. Smith, “ADRAS-J mission takes methodical first steps toward the commercial removal of space debris,” *NASA SpaceFlight.com*, Nov. 2024.
- [50] J. Rainbow, *Astroscale unlocks remaining space agency funds to de-orbit OneWeb satellite*, Jul. 2024.

- [51] T. Pultarova, “Space-junk scout captures amazing fly-around footage of discarded rocket in orbit,” *Space.com*, Nov. 2024.
- [52] Orbit Fab, *RAFTI Fueling Ports Flight Qualified and Ready for In-Space and Launch Site Fueling*, Nov. 2024.
- [53] A. H. Inc., *Historic Approach to Space Debris: Astroscale’s ADRAS-J Closes in by 50 Meters*, Nov. 2024.
- [54] Wikipedia, *ClearSpace-1*, 2025.
- [55] European Space Agency, *ESA purchases world-first debris removal mission from start-up*, Nov. 2020.
- [56] A. Parsonson, “ClearSpace Debris Removal Mission Passes Key Milestone,” *European Spaceflight*, Nov. 2024.
- [57] EVONA, “ClearSpace and Arianespace Join Forces Space Debris Removal,” *EVONA*,
- [58] R. Dipietro, “Clearspace-1: ESA and the space debris challenge,” *Space Voyaging*, Nov. 2023.
- [59] T. Hitchens, “SpaceLogistics sees potential defense market for orbital life-extension spacecraft,” *Breaking Defense*, Nov. 2022.
- [60] Intelsat, *Intelsat to Extend Life of Satellite with new Mission Extension Pod*, Nov. 2023.
- [61] J. DiMascio, “Mission Robotic Vehicle Is Poised For 2024 Launch,” *Aviation Week Network*, Nov. 2023.
- [62] NorthropGrumman, *SpaceLogistics Services*, 2023.
- [63] S. Kuper, “Northrop Grumman achieves key milestone for next-generation robotic satellite servicing vehicle,” *Space Connect*, Nov. 2025.
- [64] Orbit Fab, *News*, 2024.
- [65] C. Albon, “Orbit Fab ships first satellite refueling ports to Space Force,” *C4ISRNET*, Nov. 2024.
- [66] Orbit Fab, *Space Force Approves Orbit Fab’s RAFTI as an Accepted Refueling Interface for Military Satellites*, Nov. 2024.

- [67] Orbit Fab, *UK Pioneers Space Refuelling with £2m for Orbit Fab's Innovative Technologies*, Nov. 2024.
- [68] SatNews, “Orbit Fab delivers 1st GRIP in-space refueling nozzle following successful testing,” *SatNews*, Nov. 2024.
- [69] European Space Agency, *ESA to build first in-orbit servicing mission with D-Orbit*, Nov. 2024.
- [70] NASA, *On-orbit Servicing, Assembly, and Manufacturing 1 (OSAM-1)*, Nov. 2024.
- [71] GMI, “On-orbit Satellite Servicing Market Size, By Service, Orbit Type, Satellite Type, End Use Analysis, Share, Growth Forecast 2025 – 2034,” Tech. Rep., Nov. 2024.
- [72] Wikipedia, *Astroscale*, 2025.
- [73] Telespazio, *In Orbit Servicing*, 2024.
- [74] Thales Alenia Space, “On-orbit servicing set to shift up a gear with Space USB project,” *Thales Alenia Space*, Nov. 2024.
- [75] European Space Agency, *ESA moves ahead with In-Orbit Servicing missions*, 2023.
- [76] Precedence Research, “On-Orbit Satellite Servicing Market Size to Hit USD 11.56 Billion by 2034,” Tech. Rep., Nov. 2025.
- [77] Y. Ovadia, “Astroscale, Orbit Fab are helping refuel satellites in space, clean up space junk,” *CTech*, Jan. 2022.
- [78] S. Clark, *Two commercial satellites link up in space for first time*, Feb. 2020.
- [79] Astroscale, *Astroscale's ELSA-d Mission Successfully Completes Complex Rendezvous Operation*, May 2022.
- [80] Astroscale, *Missions*, 2023.
- [81] A. O'Leary *et al.*, *The Enhanced Economics, Incentives, and Multinational Co-operation Enabled by Refueling Architectures Centered Around Debris Clusters for Sustainable Active Debris Removal*. Sep. 2022.
- [82] OrbitFab, *Stewarding innovations in space for Earth*, 2023.

- [83] E. Spessert, Z. Burkhardt, C. Hernandez, A. O’Leary, J. Bultitude, and D. Faber, “Designing a Commercial Mission Operations Center for RPOD Missions,” Mar. 2023.
- [84] A. Jones, *NASA cancels multibillion-dollar on-orbit satellite servicing mission*, Mar. 2024.
- [85] M. Lucas-Rhimbassen, C. Santos, G. A. Long, and L. Rapp, “Conceptual model for a profitable return on investment from space debris as abiotic space resource,” in *8 TH EUROPEAN CONFERENCE FOR AERONAUTICS AND SPACE SCIENCES (EUCASS)*, Madrid: EUCASS, Jul. 2019.
- [86] P. J. Sears and K. Ho, “Impact Evaluation of In-Space Additive Manufacturing and Recycling Technologies for On-Orbit Servicing,” *Journal of Spacecraft and Rockets*, vol. 55, no. 6, pp. 1498–1508, Aug. 2018.
- [87] F. Koch, “A BUSINESS CASE FOR SPACE DEBRIS EXECUTIVE SUMMARY,” Ph.D. dissertation, Berlin, Jul. 2021.
- [88] J. Wright, “How to Recycle Space Debris and Reduce Launch Costs,” *NOW*, Mar. 2019.
- [89] T.-Y. Fung, S. S. Roy, Q. Shi, and D. A. DeLaurentis, “Space Junk Aggregation, Neutralization, In-situ Transformation, and Orbital Recycling,” in *2022 17th Annual System of Systems Engineering Conference (SOSE)*, IEEE, Jun. 2022, pp. 239–245, ISBN: 978-1-6654-9623-0.
- [90] D. Messier, “NASA Funded R&D Projects Aim to Recycle Defunct Satellites, Upper Stages,” *Parabolic Arc*, May 2022.
- [91] M. Luu and D. E. Hastings, “Valuation of On-Orbit Servicing in Proliferated Low-Earth Orbit Constellations,” in *ASCEND 2020*, ser. ASCEND, American Institute of Aeronautics and Astronautics, Nov. 2020.
- [92] B. Meng *et al.*, “The orbit deployment strategy of OOS system for refueling near-earth orbit satellites,” *Acta Astronautica*, vol. 159, pp. 486–498, 2019.
- [93] T. Sarton du Jonchay, H. Chen, O. Gunasekara, and K. Ho, “Framework for Modeling and Optimization of On-Orbit Servicing Operations Under Demand Uncertainties,” *Journal of Spacecraft and Rockets*, vol. 58, no. 4, pp. 1157–1173, Jun. 2021.
- [94] R. S. Legge, “Optimization and Valuation of Reconfigurable Satellite Constellations Under Uncertainty,” Ph.D. dissertation, Massachusetts Inst. of Technology, 2014.

[95] C. Geiman *et al.*, *A Shuttle and Depot Architecture for Reliable and Cost-Effective Refueling Operations in All Orbits*. Sep. 2022.

[96] Z. Zhao, J. Zhang, H.-y. Li, and J.-y. Zhou, “LEO cooperative multi-spacecraft refueling mission optimization considering J2 perturbation and target’s surplus propellant constraint,” *Advances in Space Research*, vol. 59, no. 1, pp. 252–262, 2017.

[97] K. Ho, H. Wang, P. A. DeTrempe, T. S. du Jonchay, and K. Tomita, “Semi-Analytical Model for Design and Analysis of On-Orbit Servicing Architecture,” *Journal of Spacecraft and Rockets*, vol. 57, no. 6, pp. 1129–1138, May 2020.

[98] O. L. de Weck, U. Scialom, and A. Siddiqi, “Optimal reconfiguration of satellite constellations with the auction algorithm,” *Acta Astronautica*, vol. 62, no. 2-3, pp. 112–130, Jan. 2008.

[99] H. W. Lee and K. Ho, “Regional Constellation Reconfiguration Problem: Integer Linear Programming Formulation and Lagrangian Heuristic Method,” *Journal of Spacecraft and Rockets*, vol. 60, no. 6, pp. 1828–1845, Jul. 2023.

[100] J. Rainbow, *Astroscale clears critical design review for OneWeb satellite removal demo*, Jun. 2025.

[101] J. Rainbow, *Astroscale gets funds for 2024 debris-removal mission*, May 2022.

[102] N. Ferraro, *Amazon wants FCC to drop five-year rule on deorbiting satellites*, Sep. 2025.

[103] C. Kuersten, “NOTICE OF EXPARTE,” Kuiper Systems LLC, Tech. Rep., Sep. 2025.

[104] S. Cornara, T. Beech, M. Bello-Mora, and A. Martinez de Aragon, “Satellite constellation launch, deployment, replacement and end-of-life strategies,” in *13th Annual AIAA/USU Conference on Small Satellites*, 1999.

[105] N. T. Redd, “Zombie Satellites Return From the Graveyard,” *IEEE Spectrum*, Jun. 2020.

[106] *Northrop Grumman Achieves First-Ever Undocking Between Two Commercial Spacecraft in Geosynchronous Orbit*, Apr. 2025.

[107] *ESA to build first in-orbit servicing mission with D-Orbit*, Oct. 2024.

[108] Stoke, “The Independent Investment Guide,” Stoke, Kent, Tech. Rep., Feb. 2024.

- [109] R. Parker, *On-orbit satellite servicing, insurance and lessons of Palapa B2 and Westar 6*, 2015.
- [110] *Last shuttle crew gets award*, Dec. 1984.
- [111] J. Uri, “35 Years Ago: STS-51A – ”The Ace Repo Company”,” *NASA History*, Nov. 2019.
- [112] J. Uri, “35 Years Ago: STS-32 Returns the Long Duration Exposure Facility,” *NASA History*, Nov. 2025.
- [113] NASA, *Long Duration Exposure Facility (LDEF)*, 2012.
- [114] eoPortal, *LDEF (Long Duration Exposure Facility)*.
- [115] A. S. Levin, “LDEF 69 Months in Space: First Post-Retrieval Symposium,” in *NASA Conference Publication 3134*, NASA, Jun. 1991.
- [116] NASA, *STS-57*, Nov. 2025.
- [117] G. D. Krebs, *Eureca 1*.
- [118] H. Hubner, P. Ferri, and W. Wimmer, “EURECA MISSION CONTROL EXPERIENCE AND MESSAGES FOR THE FUTURE,” European Space Operations Centre, Tech. Rep., Nov. 1994.
- [119] NASA, *STS-72*, Nov. 2025.
- [120] Institute of Space and Astronautical Science, *Past: SFU*, 2025.
- [121] B. Evans, “Mission for Japan: 20 Years Since STS-72 (Part 1),” *AmericaSpace*, 2015.
- [122] A. A. Boysen, T. A. Kunsa, K. J. Okseniuk, and H. R. Magill, “An Economic Case for Distributed, On-Demand Down-Mass Systems in Low Earth Orbit,” in *AIAA AVIATION FORUM AND ASCEND 2024*, ser. AIAA Aviation Forum and ASCEND co-located Conference Proceedings, American Institute of Aeronautics and Astronautics, Jul. 2024.
- [123] G. Taguchi, *System of experimental design : engineering methods to optimize quality and minimize costs* /. White Plains, N.Y: UNIPUB/Kraus International Publications, 1987, ISBN: 0527916218.
- [124] R. de Neufville and S. Scholtes, *Flexibility in Engineering Design*. The MIT Press, 2011, ISBN: 9780262303569.

- [125] T. Wang and R. de Neufville, “Real Options ”in” Projects,” in *9th Real Options Annual International Conference*, Paris, Jun. 2005.
- [126] T. Mikaelian, D. J. Nightingale, D. H. Rhodes, and D. E. Hastings, “Real Options in Enterprise Architecture: A Holistic Mapping of Mechanisms and Types for Uncertainty Management,” *IEEE Transactions on Engineering Management*, vol. 58, no. 3, pp. 457–470, 2011.
- [127] S. C. Myers, “Finance Theory and Financial Strategy,” *Interfaces*, vol. 14, no. 1, pp. 126–137, Feb. 1984.
- [128] F. Black and M. Scholes, “The Pricing of Options and Corporate Liabilities,” *Journal of Political Economy*, vol. 81, no. 3, pp. 637–654, 1973.
- [129] J. C. Cox, S. A. Ross, and M. Rubinstein, “Option pricing: A simplified approach,” *Journal of Financial Economics*, vol. 7, no. 3, pp. 229–263, 1979.
- [130] A. M. Caunhye and M.-A. Cardin, “An approach based on robust optimization and decision rules for analyzing real options in engineering systems design,” *IIE Transactions*, vol. 49, no. 8, pp. 753–767, Aug. 2017.
- [131] S. Zhang and V. Babovic, “A Real Options Approach to the Design and Architecture of Water Supply Systems Using Innovative Water Technologies Under Uncertainty,” *Journal of Hydroinformatics*, vol. 14, pp. 13–29, Jan. 2012.
- [132] C. Caputo and M.-A. Cardin, “Analyzing Real Options and Flexibility in Engineering Systems Design Using Decision Rules and Deep Reinforcement Learning,” *Journal of Mechanical Design*, vol. 144, no. 2, Sep. 2021.
- [133] M.-A. Cardin, S. Zhang, and W. J. Nuttall, “Strategic real option and flexibility analysis for nuclear power plants considering uncertainty in electricity demand and public acceptance,” *Energy Economics*, vol. 64, pp. 226–237, 2017.
- [134] H. Chen, B. Gardner, P. T. Grogan, and K. Ho, “Flexibility Management for Space Logistics Through Decision Rules,” in *ASCEND 2020*, ser. ASCEND, American Institute of Aeronautics and Astronautics, Nov. 2020.
- [135] M.-A. Cardin, M. Ranjbar-Bourani, and R. de Neufville, “Improving the Life-cycle Performance of Engineering Projects with Flexible Strategies: Example of On-Shore LNG Production Design,” *Systems Engineering*, vol. 18, no. 3, pp. 253–268, May 2015.
- [136] M.-A. Cardin, Y. Deng, and C. Sun, “Real options and flexibility analysis in design and management of one-way mobility on-demand systems using deci-

sion rules,” *Transportation Research Part C: Emerging Technologies*, vol. 84, pp. 265–287, 2017.

[137] A. Babajide, R. de Neufville, and M.-A. Cardin, “Integrated Method for Designing Valuable Flexibility in Oil Development Projects,” *SPE Projects, Facilities & Construction*, vol. 4, no. 02, pp. 3–12, May 2009.

[138] R. de Neufville, S. Scholtes, and T. Wang, “Real Options by Spreadsheet: Parking Garage Case Example,” *Journal of Infrastructure Systems*, vol. 12, no. 2, pp. 107–111, Jun. 2006.

[139] B. Hamdan, Z. Liu, K. Ho, İ. E. Büyüctahtakın, and P. Wang, “A dynamic multi-stage design framework for staged deployment optimization of highly stochastic systems,” *Structural and Multidisciplinary Optimization*, vol. 66, no. 7, p. 162, 2023.

[140] Z. Burkhardt, A. O’Leary, E. Spessert, J. Rodelo, J. Bultitude, and D. Faber, “Operational Flexibility and Asset Retasking Enabled by In-Space Refueling,” Mar. 2023.

[141] J. H. Saleh, E. Lamassoure, and D. E. Hastings, “Space Systems Flexibility Provided by On-Orbit Servicing: Part 1,” *Journal of Spacecraft and Rockets*, vol. 39, no. 4, pp. 551–560, Jul. 2002.

[142] E. Lamassoure, J. H. Saleh, and D. E. Hastings, “Space Systems Flexibility Provided by On-Orbit Servicing: Part 2,” *Journal of Spacecraft and Rockets*, vol. 39, no. 4, pp. 561–570, Jul. 2002.

[143] J. H. Saleh, “Weaving time into system architecture : new perspectives on flexibility, spacecraft design lifetime, and on-orbit servicing,” Ph.D. dissertation, MIT, 2001.

[144] R. Nilchiani and D. E. Hastings, “Measuring the Value of Flexibility in Space Systems: A Six-Element Framework,” *Systems Engineering*, vol. 10, no. 1, pp. 26–44, Mar. 2007.

[145] T. H. Matos de Carvalho and J. Kingston, “Establishing a framework to explore the Servicer-Client relationship in On-Orbit Servicing,” *Acta Astronautica*, vol. 153, pp. 109–121, Dec. 2018.

[146] Federal Aviation Administration, “Special Report: Update of the Space and Launch Insurance Industry,” US Department of Transportation, Tech. Rep., 1998.

- [147] US Department of State, *Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, Including the Moon and Other Celestial Bodies*, Jan. 1967.
- [148] A. Lawrence, *BRIEF OF PROFESSOR ANDY LAWRENCE AS AMICUS CURIAE IN SUPPORT OF VACATUR*, Aug. 2021.
- [149] P. A. Shah, Z. J. Chen, and S. A. Rowings, *SPACE EXPLORATION HOLDINGS, LLC'S FINAL RESPONSE BRIEF AS INTERVENOR FOR RESPONDENT*, Dec. 2021.
- [150] United States Government Accountability Office, “Large Constellations of Satellites: Mitigating Environmental and Other Effects,” Sep. 2022.
- [151] N. Adilov, P. Alexander, and B. Cunningham, “An Economic Analysis of Earth Orbit Pollution,” *Environmental & Resource Economics*, vol. 60, no. 1, pp. 81–98, 2015.
- [152] M. K. Macauley, “The economics of space debris: Estimating the costs and benefits of debris mitigation,” *Acta Astronautica*, vol. 115, pp. 160–164, 2015.
- [153] A. Rao, M. G. Burgess, and D. Kaffine, “Orbital-use fees could more than quadruple the value of the space industry,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 23, pp. 12 756–12 762, Jun. 2020.
- [154] S. Sinha Roy, Q. Shi, C. Guariniello, and D. A. DeLaurentis, “Analysis of the Business Case for an On-orbit Space Debris Recycling Facility,” in *AIAA SCITECH 2023 Forum*, Reston, Virginia: American Institute of Aeronautics and Astronautics, Jan. 2023, ISBN: 978-1-62410-699-6.
- [155] R. Reesman, “ASSURANCE THROUGH INSURANCE AND ON-ORBIT SERVICING,” Aerospace Corporation, Tech. Rep., Feb. 2018.
- [156] J. Foust, *Insurers brace for ViaSat-3 claim*, Jul. 2023.
- [157] L. Head, *UK Space Agency undertakes consultation on orbital liabilities, insurance, charging and sustainability*, Jan. 2024.
- [158] J. Taylor, “Tragedy of the Space Commons: A Market Mechanism Solution to the Space Debris Problem,” *Columbia Journal of Transnational Law*, vol. 50, no. 1, pp. 253–279, 2011.
- [159] C. Bullock and R. Johanson, “Policies for incentivizing orbital debris assessment and remediation,” *MIT Science Policy Review* 2, 2021.

- [160] SpaceX, *Letter to the Federal Communications Commission*, Sep. 2023.
- [161] D. Baiocchi and W. I. Welser, “Confronting Space Debris: Strategies and Warnings from Comparable Examples Including Deepwater Horizon,” *RAND*, Nov. 2010.
- [162] L. Brocca, *The Space Junk Crisis: Why We Need the ‘Polluter Pays’ Principle in Orbit*, Apr. 2025.
- [163] D. Horan, “Laws,” in *The Dialogues of Plato*, Aug. 2025, ch. Book VII, pp. 105–118.
- [164] M. R. Khan, “Polluter-Pays-Principle: The Cardinal Instrument for Addressing Climate Change,” *Laws*, vol. 5, Sep. 2015.
- [165] K. P. Green, “Polluter pays principle,” *EBSCO Knowledge Advantage*, 2024.
- [166] A. Watson, R. J. Sears, and B. Campion, *Climate Superfund Act*, May 2024.
- [167] J. Raffetto and R. Quam-Wickham, “Vermont and New York Climate Acts are First in a Wave of Likely Climate Change Cost Recovery Laws,” *Sidley*, Jun. 2024.
- [168] L. Sarnoff, “Vermont officially becomes 1st state to charge big oil for climate change damage,” *ABC News*, May 2024.
- [169] M. Sebastian, “A New Horizon: State Climate Superfund Laws,” *Proskauer 150*, Jul. 2024.
- [170] *Governor Hochul Signs Landmark Legislation Creating New Climate Superfund*, Dec. 2024.
- [171] S. B. Boxerman, J. Roffetto, R. Quam-Wickham, S. E. Wallach, Crough. Maureen M., and J. Wedeking, *New York Passes Second-in-the-Nation Climate Change Superfund Act*, Jan. 2025.
- [172] S. Bozer, “Overview and Implications of New CAFE Standards and Penalties,” Center for Automotive Research, Tech. Rep., Apr. 2022.
- [173] A. Raviv, “Congress Eliminates Corporate Average Fuel Economy (CAFE) Penalties for Passenger Cars and Light Trucks,” *Sidley*, Jul. 2025.
- [174] FSR, “EU Emission Trading System (EU ETS),” *Florence School of Regulation*, Mar. 2024.

- [175] “About the EU ETS,” *European Commission*, 2025.
- [176] “EU emissions trading system,” European Union, Tech. Rep., 2025.
- [177] EPA, *Renewable Energy Certificates (RECs)*, Mar. 2025.
- [178] *Plastic Bag Tax Successfully Reduces Waste and Supports Sustainability Efforts*, Dec. 2024.
- [179] *British Columbia: Equity in Carbon Tax*, 2025.
- [180] *The Highway Trust Fund Explained*, Feb. 2024.
- [181] “Regional Greenhouse Gas Initiative (RGGI),” *Center for Climate and Energy Solutions*, 2021.
- [182] “Investments of Proceeds,” The Regional Greenhouse Gas Initiative, Tech. Rep., Jul. 2025.
- [183] S. Jones, “Waste Disposal Surcharges Review,” *BioCycle*, Mar. 2022.
- [184] *Covered Electronic Waste Recycling Fees Guide*, 2025.
- [185] *Electronic Waste Recycling Act of 2003*, 2025.
- [186] J. Borras, *Gogoro Batteries Power 90% of Taiwan’s Electric Scooters*, 2022.
- [187] M. Heffernan, *Apple’s iPhone recycling robot shifts gears*, Jun. 2024.
- [188] Global Market Insights, *Remanufactured Automotive Parts Market Size*, Mar. 2025.
- [189] C. Balogh and N. Djordjevic, *Can the telecom industry materialize a sustainable future with Product Reuse?* Sep. 2022.
- [190] M. McNamara, *How Policy Can Advance a Circular Battery Economy*, Jun. 2023.
- [191] C. Henry and B. Berger, *Amid concerns, OneWeb gets vague about constellation’s cost*, Sep. 2018.
- [192] P. Jakob, S. Shimizu, S. Yoshikawa, and K. Ho, “Optimal Satellite Constellation Spare Strategy Using Multi-Echelon Inventory Control,” *Journal of Spacecraft and Rockets*, vol. 56, no. 5, pp. 1449–1461, May 2019.

- [193] SETS, *Xenon, Krypton, or Argon Propellants for Hall Thruster: Efficiency, Properties, Applications, and Cost*, 2024.
- [194] Iridium, *Hydrazine – Toxic for humans, but satellites love it*, 2017.
- [195] R. Jewett, *Astroscale Hopes ADRAS-J Mission Will Lay the Groundwork for Commercial Debris Removal*, 2023.
- [196] J. Rainbow, *Eutelsat expects to return to growth next year with OneWeb*, 2023.
- [197] R. Jewett, *Eutelsat Grows Sales Despite Slower OneWeb Rollout*, Sep. 2024.
- [198] Citi Global Perspectives & Solutions, “The Dawn of a New Age,” Citi, Tech. Rep., May 2022.
- [199] N. Adilov, P. Alexander, B. Cunningham, and N. Albertson, “An analysis of launch cost reductions for low Earth orbit satellites,” *Economics Bulletin*, vol. 42, no. 3, 2022.
- [200] J. Seibert, “New Glenn vs Starship – A Detailed Comparison in 2025,” *Space Insider*, 2025.
- [201] F. Geng, G. F. Dubos, and J. H. Saleh, “Spacecraft obsolescence: Modeling, value analysis, and implications for design and acquisition,” in *2016 IEEE Aerospace Conference*, 2016, pp. 1–13.
- [202] P. Selding, *OneWeb, Eutelsat call for tighter satellite-disposal regulations; 4 of OneWeb’s 634 satellites have failed in orbit*, Jun. 2023.
- [203] J. Rainbow, *Eutelsat scales back OneWeb Gen 2 upgrade plan*, Feb. 2024.
- [204] A. J. Sandler and B. Craig, “FCC Advances Its Space and Satellite Agenda,” Pillsbury, Tech. Rep., Mar. 2024.
- [205] M.-A. Cardin, Q. Xie, T. S. Ng, S. Wang, and J. Hu, “An approach for analyzing and managing flexibility in engineering systems design based on decision rules and multistage stochastic programming,” *IISE Transactions*, vol. 49, no. 1, pp. 1–12, Jan. 2017.
- [206] R. Bellman, “On the Theory of Dynamic Programming,” *Proceedings of the National Academy of Sciences*, vol. 38, no. 8, pp. 716–719, Aug. 1952.
- [207] B. Hamdan, K. Ho, and P. Wang, “Staged-Deployment Design for Resilient Expansion Planning of Large Scale Complex Systems,” in *AIAA Scitech 2019*

Forum, ser. AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, Jan. 2019.

- [208] J. Lin, “Exploring Flexible Strategies in Engineering Systems Using Screening Models,” Ph.D. dissertation, MIT, Dec. 2008.
- [209] K. Ho, “Dynamic Network Modeling for Spaceflight Logistics with Time-Expanded Networks,” Ph.D. dissertation, MIT, Boston, 2015.
- [210] T. Ishimatsu, “Generalized Multi-Commodity Network Flows: Case Studies in Space Logistics and Complex Infrastructure Systems,” Ph.D. dissertation, MIT, 2013.
- [211] T. Ishimatsu, O. L. de Weck, J. A. Hoffman, Y. Ohkami, and R. Shishko, “Generalized Multicommodity Network Flow Model for the Earth–Moon–Mars Logistics System,” *Journal of Spacecraft and Rockets*, vol. 53, no. 1, pp. 25–38, Nov. 2015.
- [212] K. Ho, O. L. de Weck, J. A. Hoffman, and R. Shishko, “Dynamic modeling and optimization for space logistics using time-expanded networks,” *Acta Astronautica*, vol. 105, no. 2, pp. 428–443, 2014.
- [213] K. Ho, O. de Weck, J. Hoffman, and R. Shishko, “Campaign-level dynamic network modelling for Spaceflight logistics for the flexible path concept,” *Acta Astronautica*, vol. 123, Mar. 2016.
- [214] A. R. Collins, “Development of Cislunar Space Logistics Networks for Satellite Constellation Support Using Event-Driven Generalized MultiCommodity Network Flows,” Ph.D. dissertation, Air Force Institute of Technology, Nov. 2019.
- [215] H. Chen and K. Ho, “Integrated Space Logistics Mission Planning and Spacecraft Design with Mixed-Integer Nonlinear Programming,” *Journal of Spacecraft and Rockets*, vol. 55, pp. 1–17, Oct. 2017.
- [216] H. Chen, H. W. Lee, and K. Ho, “Space Transportation System and Mission Planning for Regular Interplanetary Missions,” *Journal of Spacecraft and Rockets*, vol. 56, no. 1, pp. 12–20, Nov. 2018.
- [217] H. Chen, T. Sarton du Jonchay, L. Hou, and K. Ho, “Integrated in-situ resource utilization system design and logistics for Mars exploration,” *Acta Astronautica*, vol. 170, pp. 80–92, 2020.

- [218] T. Sarton Du Jonchay, H. Chen, M. Isaji, Y. Shimane, and K. Ho, “On-Orbit Servicing Optimization Framework with High- and Low-Thrust Propulsion Tradeoff,” *Journal of Spacecraft and Rockets*, vol. 59, pp. 1–16, Jul. 2021.
- [219] T. Sarton du Jonchay, Y. Shimane, M. Isaji, H. Chen, and K. Ho, “ON-ORBIT SERVICING LOGISTICS FRAMEWORK GENERALIZED TO THE MULTI-ORBIT CASE,” in *AAS/AIAA Astrodynamics Specialist Conference*, Big Sky, Aug. 2021.
- [220] M. Silver and O. de Weck, “Time-Expanded Decision Network: A New Framework for Designing Evolvable Complex Systems,” *Systems Engineering*, vol. 10, pp. 167–188, Jun. 2007.
- [221] A. Stockley, *The Rise of Reusable Rockets: Transforming the Economics of Space Travel*, Apr. 2023.
- [222] S. Erwin, *DARPA’s satellite servicing robot to get another shot*, Jul. 2019.
- [223] P. Gannon, *Another MEV Rescue Mission*, Sep. 2020.
- [224] J. Locke, T. J. Colvin, L. Ratliff, A. Abdul-Hamid, and C. Samples, “Cost and Benefit Analysis of Mitigating, Tracking, and Remediating Orbital Debris,” NASA, Tech. Rep., 2024.
- [225] D. Vallado, *Fundamentals of Astrodynamics and Applications*, 5th. Microcosm Press, 2022.
- [226] S. R. Oleson, “An Analytical Optimization of Electric Propulsion Orbit Transfer Vehicles,” NASA Contractor Report, Brook Park, Ohio, Tech. Rep., 1993.
- [227] J. I. M. Volosin, *Psyche, Falcon Heavy*, Oct. 2023.
- [228] P. Beck, *Electron, Payload User Guide*, Nov. 2022.
- [229] Stoke, *Nova: Engineered for full and rapid reuse*, 2025.
- [230] J. Taylor, “Press Kit, SpaceX CRS-6 Mission,” SpaceX, Tech. Rep., Apr. 2015.
- [231] P. Beck and M. Darley, “Return To Sender: Lessons Learned From Rocket Lab’s First Recovery Mission,” in *Small Satellite Conference*, Aug. 2021.
- [232] L. Wiener, *Articulated Locomotives*, 1930.
- [233] Z. Bobbitt, *Welch’s t-test: When to Use it + Examples*, Dec. 2020.

- [234] P. Virtanen *et al.*, “SciPy 1.0: fundamental algorithms for scientific computing in Python,” *Nature Methods*, vol. 17, no. 3, pp. 261–272, 2020.
- [235] S. Lomuscio, *Bootstrap Estimates of Confidence Intervals*, Mar. 2023.
- [236] “Rocket Lab Electron,” *Wikipedia*,
- [237] “Rocket Lab Neutron,” *Wikipedia*,
- [238] “Terran R,” *Wikipedia*,

VITA

Madeline Bowne earned her Master of Science in Aerospace Engineering from the Georgia Institute of Technology in 2022. She holds a Bachelor of Science in Mechanical Engineering with a minor in English from Rutgers University, where she was actively involved in the Formula SAE racing team as Vice President and later as Testing and Design Lead. With a strong foundation in mechanical, aerospace, and systems engineering, her graduate research focused on combining advanced design and simulation expertise for improved decision-making. She presented research related to her dissertation at the 2025 AIAA SciTech and ASCEND conferences, and delivered a Lightning Talk at SGx 2025. Madeline has completed seven internships across the aerospace industry, including at Gravitics, Relativity Space, Redwire, Lockheed Martin, and Northrop Grumman. A 2021 Matthew Isakowitz Fellow, she is passionate about the commercial space sector and space sustainability.