Problem Chosen	2026 MCM/ICM Summary Sheet	Team Control Number
B		2618656

Energy as Currency: Optimizing Hybrid Transport for the 2050 Lunar Frontier

Summary

To address the logistical challenge of transporting 100 million metric tons of material to establish a 100,000-person Moon colony by 2050, we develop a Universal Energy-Equivalent and Temporal Co-ordination Model and a Life-Support Logistics and Stochastic Water Balance Model. These models evaluate the trade-offs between the Space Elevator System and traditional rocket launches across energy, time, and environmental dimensions.

Firstly, we establish a **Universal Energy-Equivalent (UEE)** metric to facilitate a thermodynamically consistent comparison between chemical rockets and electric space elevators. We introduce a **Time-Opportunity Parameter** (λ) to transform the energy-time trade-off into a single optimization objective. To ensure robustness, we incorporate **CVaR-style risk adjustments** and **Monte Carlo simulations** to account for system failures, tether swaying, and operational downtime.

For task 1, we compare three delivery scenarios. We find that while a **Rocket-Only** approach offers the shortest initial timeline, it is energetically prohibitive. The **Elevator-Only** scenario requires 186 years but consumes the least energy. The **Balanced Hybrid** scenario (139 years) emerges as a strategic compromise, balancing construction velocity with resource efficiency.

For task 2, we evaluate system reliability under non-ideal conditions. Our results indicate that the space elevator's throughput is highly sensitive to tether stability. However, even with a **15% downtime margin**, the elevator remains the superior long-term infrastructure compared to the high failure-cost risks of mass rocket launches.

For task 3, we develop a **Tiered Water Logistics Model** based on three comfort levels. Using sensitivity analysis, we identify **recycling efficiency** (η) as the dominant lever; a 1% drop in η increases annual supply needs by 9.6%. We conclude that the space elevator can comfortably support a Luxury tier, occupying 69.68% of its annual capacity.

For task 4, we extend the model into an **Environmental Single-Objective Framework**. By quantifying CO_2 emissions and stratospheric H_2O injection, we find the **Elevator-Only** plan reduces carbon footprints by 93.5% compared to rockets. By synthesizing temporal efficiency, energy consumption, and environmental externalities, the model identifies the optimal transport strategy for lunar colonization that prioritizes the preservation of Earth's ecological integrity.

Finally, we recommend a **Tiered Strategy**: beginning with Survival-tier logistics to secure the colony, then transitioning to a Comfort-tier elevator-based operation to achieve long-term sustainability.

Keywords: Energy-based cost proxy; Stochastic Risk Analysis ; Multi-objective Optimization: Earth–Moon logistics: Space Elevator

Contents

1 Introduction

1.1 Background

As human space exploration shifts from short-term landings to permanent settlement, Earth's finite resources and fragile ecosystems motivate the development of scalable off-world habitats. Yet traditional chemical-propulsion rockets face persistent bottlenecks—limited payload ratios, high marginal launch costs, and non-negligible environmental impacts—making it difficult to sustain the mass logistics required for constructing and operating a 100,000-person Moon Colony. In this context, the Space Elevator System has been proposed as a transformative infrastructure [6] that can provide routine, energy-efficient access to space when paired with lunar transfer, but it also introduces new reliability and operational uncertainties.

1.2 Restatement of the Problem

We develop a quantitative model to estimate the cost proxy (energy) and timeline of transporting materials to build a 100,000-person Moon Colony starting in 2050, and to support a defensible recommendation. Specifically, we:

- Compare three delivery scenarios: elevators-only, rockets-only (selected sites), and a hybrid portfolio;
- Test robustness under non-ideal operations (e.g., tether dynamics, failures, downtime, and launch interruptions);
- Quantify one-year water resupply needs after habitation and map them to the delivery model;
- Evaluate and mitigate environmental impacts across scenarios;
- Deliver a one-page recommendation letter to the MCM Agency.

1.3 Our Work

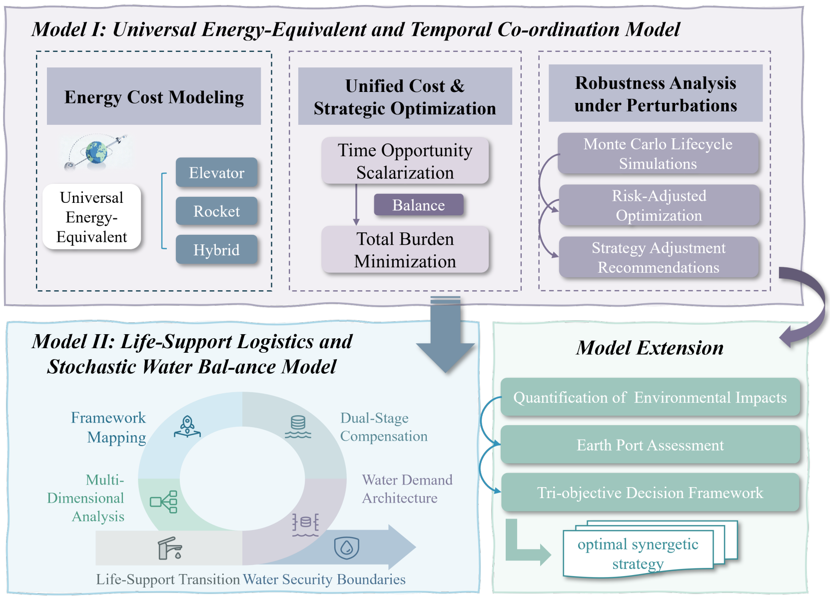

Figure 1: Our Work

2 Assumptions and Justifications

To ensure model tractability while maintaining physical realism, we establish the following assumptions:

- **Assumption 1:** Rocket payload capacity is assumed to be standardized at 100–150 metric tons per mission.
  Justification: Constant capacity avoids the volatility of payload metrics, facilitating a streamlined and efficient analysis of the total energy-cost required for construction.
- **Assumption 2:** Rocket propulsion systems are assumed to utilize liquid oxygen and methane (LOX/CH_4) as propellant.
  Justification: Current reusable heavy-lift standards, such as the SpaceX Starship, utilize LOX/CH_4 for its high energy density and minimal carbon buildup, representing future interplanetary transport norms.
- **Assumption 3:** Atmospheric drag and energy losses due to vehicle structural mass are neglected during rocket transit.
  Justification: These losses vary significantly per mission. Their exclusion reduces model complexity without compromising the macro-level accuracy of the energy expenditure framework.
- **Assumption 4:** Capital expenditures, including R&D, manufacturing, and labor costs, are excluded from the assessment.
  Justification: This study focuses on the operational application phase of the logistics chain. Furthermore, such costs are often proprietary and difficult to estimate for future technologies.
- **Assumption 5:** Industrial water demand is assumed to be excluded from the water replenishment model.
  Justification: Lunar industrial scale data remains unavailable in current literature. Such demand is assumed to be integrated into the primary construction material transport modeled in Task 1.

3 Notations

Symbols	Description	Unit
m_{fuel}	Total mass of the chemical fuel required for a mission	metric tons
m_p	Mass of the effective payload being transported	metric tons
R	Mass ratio of initial mass to final mass	-
E	Total energy expenditure (Universal Energy-Equivalent)	Joules
h	Length of the space elevator tether	km
t	Total duration required for the construction timeline	years
p	Probability of rocket launch failure	-
W	Total water replenishment demand for the colony	metric tons

4 Model I: Universal Energy-Equivalent and Temporal Coordination Model

4.1 Model Overview

To evaluate Earth–Moon logistics, we propose the Universal Energy-Equivalent and Temporal Coordination Model (UETCM). We use Universal Energy-Equivalent (UEE, J) as a physics-based cost proxy to compare chemical rockets and the space elevator on a common basis.

UETCM maps payload mass to energy using the Tsiolkovsky equation and gravitational mechanics, and reconciles the energy–time trade-off by introducing a time opportunity cost parameter λ to form a single objective J.

We further incorporate non-ideal operations via a stochastic extension with Monte Carlo simulation and Conditional Value-at-Risk (CVaR), yielding risk-aware strategies with explicit robustness margins.

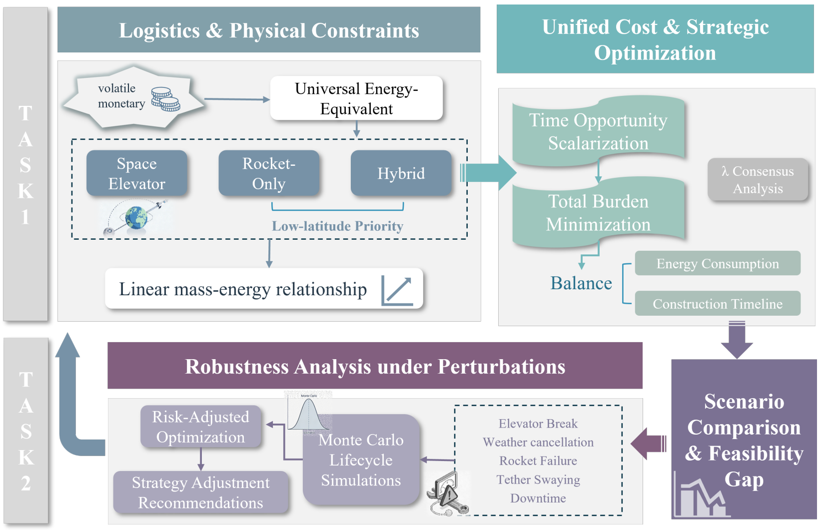

Figure 2: Flow Chart of Model I

4.2 Energy-Equivalent as a Cost Proxy

Monetary costs are ill-conditioned for an Earth–Moon logistics program spanning many decades due to inflation, technology shocks, and geopolitical uncertainty [5]. We therefore use *energy consumption* as a stable, physics-grounded proxy for marginal cost, enabling a consistent comparison between rockets and the Space Elevator System.

4.2.1 Cost Convergence Analysis

The operational cost per mission, $Cost(x)$, is modeled as the sum of hardware amortization, energy expenditure, and maintenance expenses:

$$Cost(x) = \frac{C_{hardware}}{x} + C_{energy} + C_{operations} \tag{1}$$

where x is the number of reuse cycles. As reuse increases in a mature post-2050 transport regime, $\frac{C_{hardware}}{x}$ diminishes and operations become more automated, leaving C_{energy} as the dominant term. Thus, differences in energy efficiency largely determine the marginal cost of sustained high-throughput transport.

4.2.2 Empirical Industry Validation

Evidence from mature terrestrial transport supports this modeling choice: in commercial aviation, fuel remains a persistent and material share (15–25%) of operating expenses, indicating that energy is a primary driver once hardware is amortized and operations are standardized [9].

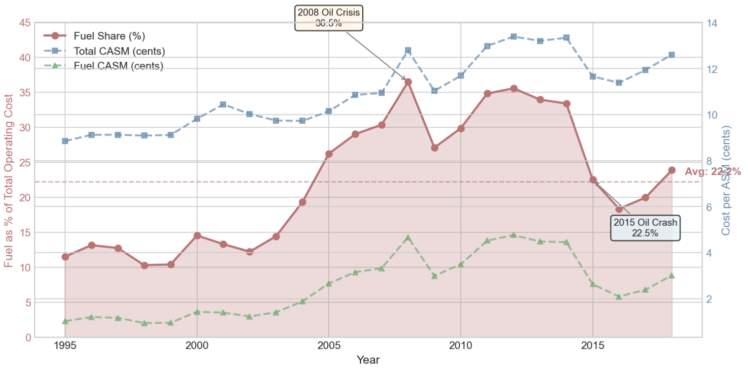

Data Source: MIT Airline Data Project / US DOT Form 41

Figure 3: Trend of Fuel Cost as a Percentage of Total Operating Expenses

In summary, we adopt Universal Energy-Equivalent as the primary metric to avoid long-horizon monetary ambiguity and to compare transport architectures on a common thermodynamic baseline (i.e., overcoming gravitational wells and imparting the required kinetic energy).

4.3 Ideal Energy Cost Modeling

Rocket Momentum Dynamics:

For traditional rocket transport, energy demand is dominated by propellant required to deliver a payload to the Earth–Moon transfer trajectory. Neglecting air resistance, the Tsiolkovsky rocket equation gives the mass ratio R:

$$R = \frac{m_0}{m_f} = e^{\Delta V / v_e}$$

where v_e is the effective exhaust velocity and ΔV (e.g., for Trans-Lunar Injection) is treated as a mission constant. Accounting for structural mass via a structural coefficient α (structure-to-fuel ratio), the fuel mass scales linearly with payload:

$$m_{fuel} = \underbrace{\frac{R-1}{1-\alpha(R-1)}}_{k} \cdot m_p$$

which defines the proportional coefficient k. Using the idealized chemical energy release proxy and under fixed vehicle parameters,rocket energy scales linearly with payload mass. For completeness, we also express the total energy requirement as the sum of gravitational and transition terms:

$$\Delta E_{total} = m_p \cdot \left[\left(\frac{GM_E}{R_E} - \frac{GM_E}{d_{EM}} \right) - \frac{GM_M}{2r_M} \right]$$

which is likewise proportional to m_p.

Space Elevator Mechanics:

Compared with direct rockets, the Space Elevator System adds an Earth-surface-to-apex segment along the tether before the lunar-transfer rocket leg (already covered above). We model only the elevator segment energy change.

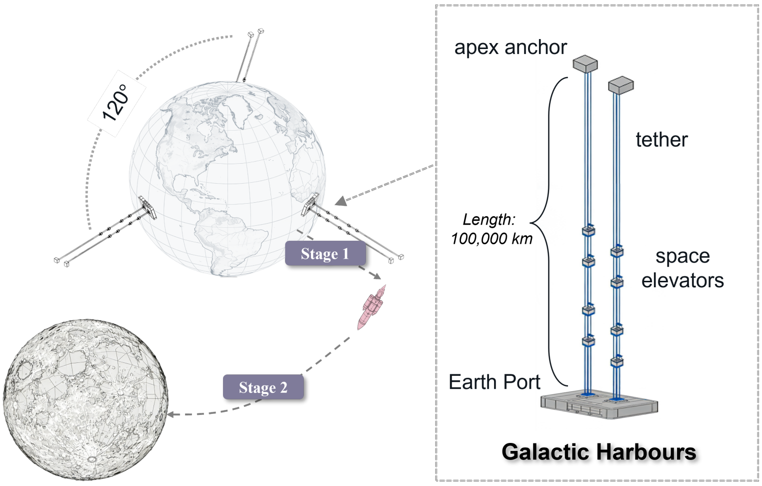

Figure 4: Schematic of the Space Elevator Infrastructure and Earth-Moon Transport Phases

From the Earth Port to the Apex Anchor, the energy change includes gravitational potential gain and rotation-related kinetic terms:

$$E_{elevator} = m_p \left[\left(\frac{GM_E}{R_E} - \frac{GM_E}{r_0} \right) + \frac{1}{2}\omega^2 r_0^2 \right]$$

where r_0 is the distance from the Apex Anchor to Earth's center and ω is Earth's angular velocity; thus $E_{elevator}$ also scales linearly with m_p.

4.4 Ideal Timeline and Logistic Efficiency Modeling

After resolving the physical energy efficiency, the model determines the construction timeline. The complexity arises from the fact that 100 million metric tons of material cannot be delivered instantaneously; the progress is constrained by the geographical distribution of launch sites, frequency limits, and the inherent capacities of the space elevator and rocket systems.

1. Transport Progress Formulas:

To quantify the construction timeline, the model defines $M_{total} = 10^8$ metric tons as the total demand. Based on the three scenarios provided, the completion time t depends on the annual cumulative capacity C:

- Space elevator only (Scenario a): $t_a = \frac{M_{total}}{C_e}$
- Rockets only (Scenario b): $t_b = \frac{M_{total}}{C_r}$
- Hybrid transport (Scenario c): $t_c = \frac{M_{total}}{w_e C_e + (1-w_e) C_r}$

We interpret the problem-stated throughput of 179,000 t/yr as the annual capacity of a single Galactic Harbour; therefore, the three-harbour Space Elevator System has a combined capacity of

$$C_e = 3 \times 179,000 \text{ t/yr}$$

. And C_r is the integrated annual capacity of the rocket system. Shortening the timeline requires increasing annual capacity. By adjusting the hybrid weights λ and μ, we seek a dynamic balance between construction speed and energy cost.

2. Spatial and Geographical Optimization

The primary obstacles to construction progress are geographical constraints and frequency limits, which define the upper bound of C_r.

- **Frequency constraints:**

To estimate the global launch capacity ceiling, we apply a Richards growth model to historical launch data. As shown in Figure 5, the model predicts a theoretical saturation limit of $K = 4298$ annual launches. However, accounting for practical constraints such as maintenance cycles and meteorological windows, we adopt a conservative operational cadence of one launch per day per site (365 annually). This calibrated cap serves as the primary constraint for determining the minimum construction duration.

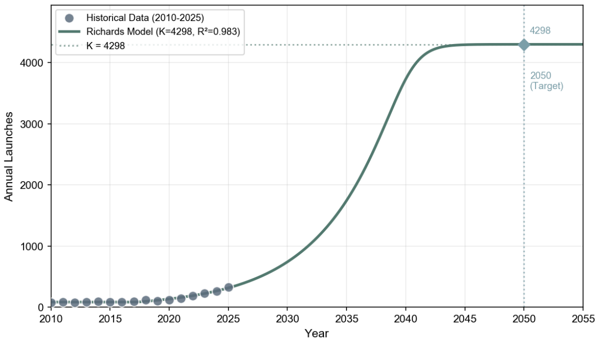

Figure 5: Richards growth model fitting and projected capacity limit K

- **Latitude-based efficiency correction:**

Earth's rotation provides rockets with a tangential initial velocity v_{0i}. The energy gain varies by latitude as follows:

$$v_{0i} = v_{0E} \cdot \cos\theta_i$$

As illustrated in Figure 6, higher latitudes result in lower initial velocities, necessitating more chemical fuel. As latitude increases, the non-linear loss of tangential velocity causes the fuel-to-payload ratio to increase from 31.8 at equatorial sites to 34.0 at high-latitude sites.

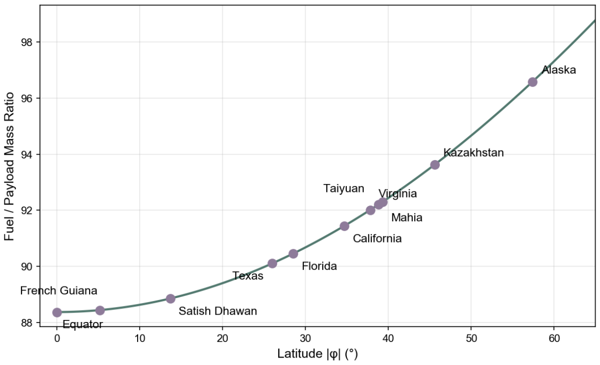

Figure 6: Rotation Velocity and Loss vs Latitude

Consequently, the model implements a low-latitude priority principle: low-latitude sites such as Kourou are prioritized to minimize global energy consumption when the number of launches has not reached the limit. This ensures that for a given timeline, the efficiency of material delivery is maximized.

4.5 Unified Cost Optimization via Time Opportunity Scaling

To select an actionable transport strategy, we scalarize the two objectives—energy consumption (E) and construction duration (T)—using a time opportunity cost parameter λ.

4.5.1 Reformulation via Opportunity Cost

We define the Total System Burden as:

$$J(T) = E(T) + \lambda T, \qquad \left.\frac{dE}{dT}\right|_{T=T^*} = -\lambda \tag{2}$$

Here λ (PJ/year) is an energy-equivalent penalty per additional year, capturing the value of schedule acceleration (e.g., earlier utilization and reduced long-horizon overhead). The optimal duration T^* is reached when the marginal energy saving rate equals λ.

4.5.2 Calibration of Cost Parameter λ

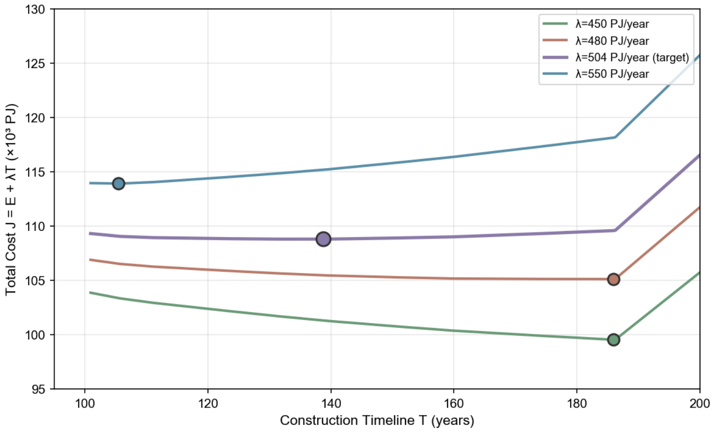

Figure 7: Sensitivity analysis of optimal timeline T^* versus time opportunity cost λ

Rather than choosing λ subjectively, we cross-check it using four independent heuristics (normalized trade-off scaling, marginal-slope matching, phase-boundary location, and feasible-region geometry). These estimates converge near $\lambda \approx 504$ PJ/year, so we use $\lambda = 504$ PJ/year as the baseline and treat the spread as sensitivity.

4.6 Results of Task 1

Through numerical simulation of the 100 million metric ton material transport mission, this study elucidates the intrinsic correlation between construction progress and resource consumption under diverse technical constraints. The findings demonstrate that the optimal transport strategy is not a simple superposition of individual methods but a dynamic equilibrium based on the trade-off between physical energy efficiency and temporal costs.

1. Comparison of Baseline Scenarios

Based on the model outputs, Table 1 summarizes the core indicators for the three baseline scenarios.

Table 1: Multi-indicator Comparison of Earth-Moon Material Transport Scenarios

Scenario Type	Min Duration (a)	Total Energy (PJ)	Unit Energy (GJ/t)
Elevator-Only (a)	186.2	15,720	157.2
Rocket-Only (b)	219.2	50,609	506.1
Hybrid (c)	100.7	31,537	315.4

Integrating Table 1 and Figure 8, a clear duration-energy trade-off emerges. Rocket-only is limited to 219.2 years with peak energy demand, while elevator-only provides the lowest energy footprint but requires 186.2 years. The hybrid scenario overcomes these bottlenecks, compressing the timeline to 100.7 years. It is the sole solution for the 100–186 year window and converges to elevator efficiency as duration increases. Ultimately, hybrid transport is indispensable for speed, while the space elevator represents the optimal long-term balance of time and energy.

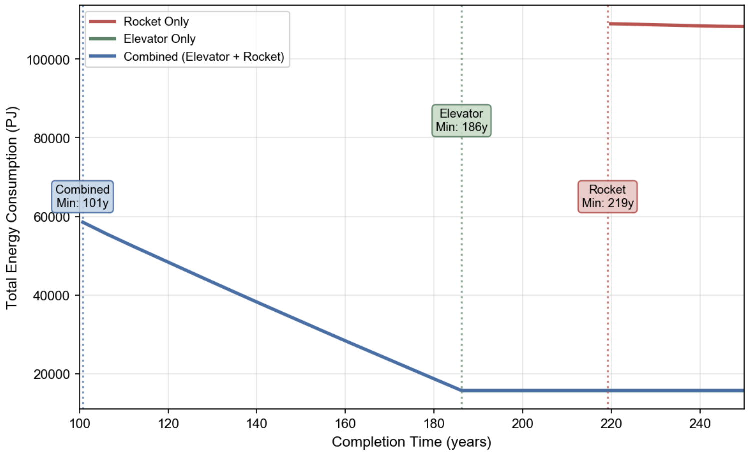

Figure 8: Three Transport Scenarios Feasibility Comparison

2. Optimal Strategy Selection and Allocation To determine the optimal operating point, we analyze the behavior of the total cost function $J(T) = E(T) + \lambda T$. Based on the consensus value of $\lambda = 504$ PJ/year, three strategic configurations are evaluated within the feasible optimization space (100.7–186 years), as detailed in Table 2.

Table 2: Core Metrics of Typical Optimization Strategies for Scenario C

Strategy Type	Duration (a)	Total Energy (PJ)	Elevator Share	Energy Saving
Strategy A (Cost-Prioritized)	186.0	15,826	99.9%	85.5%
Strategy B (Time-Prioritized)	101.0	58,391	54.2%	46.4%
Strategy C (Balanced)	139.0	**38,729**	74.6%	**64.4%**

These strategies represent key operational configurations: Strategy A minimizes physical energy via full-load elevator operation; Strategy B achieves the shortest timeline through maximum energy input; and Strategy C represents the global minimum of the total cost function J, balancing construction efficiency with energy expenditure.

1. **Marginal Energy Saving**: As shown in Figure **??**, the marginal saving curve exhibits a step-wise diminishing trend. Between years 101 and 139, each additional year allocated to the timeline reduces energy demand by approximately 210 PJ.
2. **Optimal Operating Point**: The 139-year duration is identified as the optimal point where the marginal reduction in energy expenditure equals the marginal time opportunity cost λ. At this point, the total system burden J is minimized.

3. Sensitivity and Parameter Stability

To assess the robustness of our model, we perform sensitivity analysis on five key parameters. Figure 9 summarizes the results.

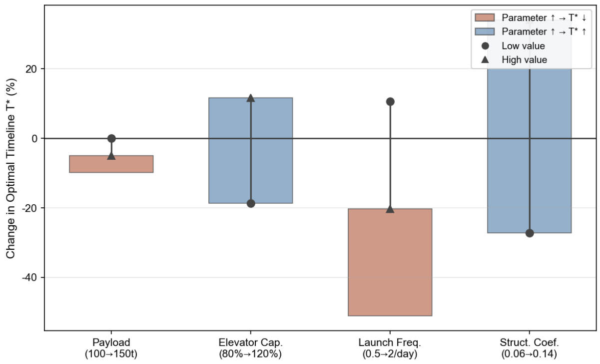

Figure 9: Sensitivity Analysis of Key Parameters for UETCM

- **Rocket Payload Capacity (100–150 tons):** As specified in the problem statement, we analyze the full payload range. Results show that higher capacity reduces minimum timeline from 111 to 93 years, but has limited impact on the optimal hybrid solution at $\lambda = 504$.

- **Elevator Capacity (±20%):** Elevator capacity emerges as the most sensitive parameter. A 20% reduction extends the elevator-only timeline from 186 to 233 years, fundamentally altering the hybrid strategy's feasible region.
- **Launch Frequency:** Doubling launch frequency from 1 to 2 per day reduces minimum timeline from 101 to 69 years, enabling more aggressive time-optimized strategies.
- **Engine Technology:** Engine Isp primarily affects energy consumption rather than timeline. LOX/Hydrogen (Isp=450s) reduces fuel ratio by 64% compared to LOX/Methane, but practical considerations favor the latter.

Table 3: Sensitivity Analysis Summary

Parameter	Range	ΔT^*	ΔE^*
Payload Capacity	100–150 t	-7 y (-5.0%)	$+3246$ PJ ($+8.4\%$)
Elevator Capacity	$\pm 20\%$	$+42$ y ($+30.2\%$)	-47902 PJ (-61.7%)
Launch Frequency	0.5–2/day	-43 y (-30.9%)	$+20867$ PJ ($+53.8\%$)
Structural Coef. α	0.06–0.14	$+85$ y ($+61.4\%$)	-28149 PJ (-72.6%)

4.7 Stochastic Risk and Robustness Analysis

Task 2 asks a fundamentally different question from Task 1: not merely how much performance degrades, but how our recommended strategy should change in response to system uncertainties. We achieve this through a three-stage methodology: quantifying perturbation impacts, identifying risk-adjusted decision boundaries, and finally deriving modified optimal strategies with robust safety margins.

(1) Stochastic Perturbation Modeling

Given the unprecedented nature of the Space Elevator System, we acknowledge significant epistemic uncertainty in failure parameters. We adopt a conservative approach by establishing parameter ranges rather than point estimates:

Table 4: Perturbation Parameter Ranges and Sources

Parameter	**Baseline**	**Range**	**Basis**
Tether swaying ($\Delta\theta$)	0°	$[-0.5^\circ, 0.5^\circ]$	Coriolis-induced oscillations during release
Rocket failure rate	0.75%	[0.3%, 1.5%]	Falcon 9 historical [3] + technology maturation
Elevator breaks/year	2	[1, 4]	IAA Space Elevator Assessment [7]
Downtime per break	14 days	[7, 30] days	Scaled from offshore platform repairs
Weather cancellation	10%	[5%, 20%]	Site-specific meteorological data

(2) Risk-Adjusted Optimization

Under uncertainty, the deterministic objective function must be extended to incorporate risk preferences. We reformulate the optimization as:

$$\min_{T} \quad J_{robust}(T) = \mathbb{E}[E(T)] + \lambda \cdot \mathbb{E}[T] + \gamma \cdot \text{CVaR}_{\alpha}(T) \tag{3}$$

where $\mathrm{CVaR}_{\alpha}(T)$ denotes the Conditional Value-at-Risk at confidence level α (typically 95%), representing the expected completion time in the worst $1-\alpha$ scenarios [5]. The risk aversion coefficient γ (PJ/year) quantifies the MCM Agency's willingness to pay for reduced schedule uncertainty.

(3) Monte Carlo Lifecycle Assessment

Due to the highly non-linear interactions between these stochastic variables, we utilized a Monte Carlo algorithm to execute 10,000 simulations of the construction cycle.

- **Simulation Logic**: For each simulated day, the algorithm samples the system state (Normal/Failure/Weather). In the event of interference, payload distribution is adjusted and incomplete tasks are rescheduled.
- **Distribution Analysis**: As shown in Figure 10, the results exhibit a significant right-skewed distribution. This implies that real-world timelines are prone to a long tail of cumulative delays—the probability of massive delays is low but the impact is profound, serving as a critical basis for risk assessment.

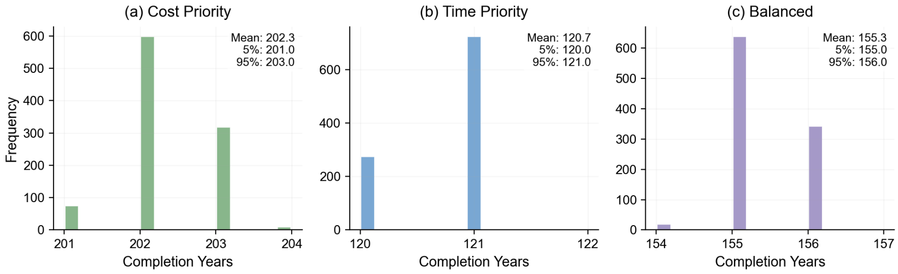

Figure 10: Completion Time Distribution Under Perturbations

4.8 Results of Task 2

Idealized models define the theoretical limits, but stochastic variables such as tether swaying, equipment failure, and weather interference cause performance to deviate. We evaluate system performance under disturbance based on 10,000 Monte Carlo iterations.

1. Assessment of Performance Degradation

Table 5 compares the deterministic solutions with stochastic means, quantifying the costs of system perturbations.

Table 5: Comparison of Core Metrics Under Perfect and Non-Perfect Conditions

Strategy Type	Dimension	Perfect Condition	Non-Perfect (Mean ± Std)	Change Rate
Strategy A	Time (a)	186.2	202.3 ± 0.6	+8.6%
	Energy (PJ)	15,720	15,738 ± 0	+0.1%
Strategy B	Time (a)	100.7	120.7 ± 0.4	+19.9%
	Energy (PJ)	31,537	30,217 ± 30	-4.2%
Strategy C	Time (a)	139.0	155.3 ± 0.5	+11.7%
	Energy (PJ)	**38,729**	**38,285 ± 45**	-1.2%

2. Resilient Strategic Adjustments

The Monte Carlo results reveal that perturbations do not merely shift performance metrics—they fundamentally alter the risk-return profile of each strategy:

(1) Optimal Timeline Adjustment

Under perfect conditions, the balanced strategy yields $T^* = 139$ years. However, incorporating the 95th percentile delay, the effective planning horizon becomes:

$$T^*_{adjusted} = T^* + \Delta T_{buffer} = 139 + 16.3 \times 1.15 \approx 158 \text{ years} \tag{4}$$

where the 15% margin accounts for tail-risk scenarios. Recommendation: MCM Agency should plan for a 155-160 year timeline to achieve 95% confidence in completion.

(2) Hybrid Ratio Adjustment

Given that elevator breaks dominate schedule risk (correlation = 0.836), the optimal elevator share should be reduced to enhance redundancy:

$$\text{Elevator share}_{adjusted} = 74.6\% - \Delta_{redundancy} \approx 65\% - 70\% \tag{5}$$

This increases rocket utilization as a "backup channel," accepting a 5-8% energy penalty in exchange for improved schedule resilience.

(3) Strategic Reserve Capacity

We recommend maintaining a 10% reserve margin in annual transport capacity:

$$C_{reserve} = 0.1 \times (\lambda C_e + \mu C_r) \tag{6}$$

This reserve can absorb localized disruptions without triggering cascading delays.

3. Evolution of Key Indicators

Stochastic disturbances transform deterministic values into probability distributions, but the strategic ranking remains robust.

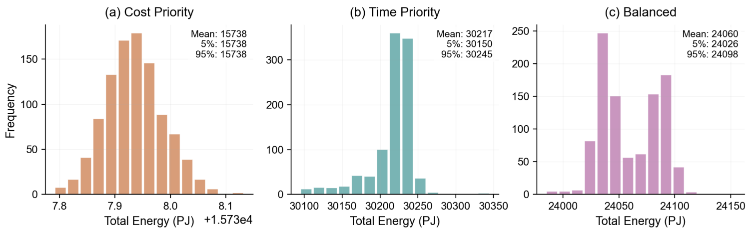

Figure 11: Energy Distribution

- **Energy Stability**: As illustrated in Figure 11, energy consumption follows a narrow-peak normal distribution with a standard deviation of less than 0.2%. This confirms that tether swaying does not cause energy collapse. The slight mean shift in Strategies B and C is due to payload loss from rocket failures offsetting energy increments from swaying.
- **Temporal Right-Skewed Characteristics**: Median durations delay by 8.6% to 19.9%. However, the duration intervals for each strategy remain strictly separated, ensuring that the fundamental decision logic—Strategy A for cost and Strategy B for speed—holds true even under perturbation.

5 Model II: Life-Support Logistics and Stochastic Water Balance Model

5.1 Model Overview

As the Moon Colony transitions from the construction phase to the operational phase, the logistical focus shifts from structural materials to life-support supplies. We develop the Life-Support Logistics and Stochastic Water Balance Model (LSL-SWBM) to quantify the water security boundaries of the settlement during its first year of operation. This model accounts for multi-level demand functions based on psychological comfort and incorporates a normal approximation to address stochastic medical emergency needs. By mapping these water requirements onto the transport framework established in Model I, we evaluate the additional pressure exerted by different comfort factors on the Earth-Moon logistics chain.

5.2 Water Demand Architecture

In an isolated lunar ecosystem, water consumption is primarily sustained by a recycling system [4]. However, logistical replenishment must compensate for system losses and sudden medical surges. This model assumes that lunar water use is restricted to domestic and medical emergency purposes, excluding industrial use, to define the core demand framework.

5.2.1 Domestic Water Evolution

Within the colony, domestic water consists of survival and hygiene components. The daily demand is formulated as:

$$W_r = N \times (w_s + 0.4\kappa) \tag{7}$$

where N is the population, w_s represents the survival baseline, and κ is the comfort factor. The following table summarizes the water standards across different demand tiers [1][8]:

Table 6: Water Demand Standards Across Different Comfort Tiers

Demand Tier	**κ Value**	**Daily Use per Capita**	**Description**
Survival Standard	1	2.9 L	Minimum survival threshold
Comfort Standard	50	22.5 L	Moderate domestic comfort
Luxury Standard	250	102.5 L	Equivalent to Earth-like usage

5.2.2 Medical Emergency Demand Modeling

Given the large population , the daily number of patients X (assuming a daily incidence rate $p = 2\%$) follows a binomial distribution, which is accurately approximated by a normal distribution $X \sim N(Np, Np(1-p))$. To ensure medical safety under extreme conditions, the model adopts the peak demand at a 99% confidence level as the daily reserve target:

$$W_{medical} = (E(X) + Z_{0.99} \cdot \sigma) \times 5 \text{ kg} \tag{8}$$

This indicator remains stable across different domestic comfort scenarios, ensuring the robustness of the medical support system.

5.3 Replenishment and Buffer Strategies

Unlike construction materials, water is highly recyclable. Given the maturity of water recycling technology, the cumulative rate of daily demand after the initial transport is relatively slow. Since a daily delivery schedule would be prohibitively expensive, we adopt a monthly supply mode based on a dual-stage replenishment logic.

5.3.1 Initial Month Filling

During the first month, recycled water from the previous month is unavailable. Thus, the initial supply must satisfy two criteria: providing domestic water for 30 days and maximizing medical reserves for potential surges. The initial transport volume is defined as:

$$W_{initial} = (W_r + W_{mi}) \cdot T \tag{9}$$

where $T = 30$ days and W_{mi} is the daily emergency medical supply at a 99% confidence level. This strategy establishes the system's circulating base while creating a 30-day emergency buffer.

5.3.2 Monthly Routine Compensation

In subsequent months, domestic demand is met through a combination of Earth-based replenishment and water recycling. Simultaneously, medical reserves need only cover the mean incidence rate due to the confidence buffer established initially. With a recycling efficiency η (set at 0.9), the routine monthly supply model is:

$$W_{routine} = (W_r(1-\eta) + W_m) \cdot T \tag{10}$$

This model precisely offsets recycling losses and daily medical consumption to maintain a dynamic water balance.

5.4 Results of Task 3

To evaluate the operational sustainability of the Moon Colony, we perform a quantitative decomposition of the water replenishment mission across temporal and energetic dimensions. By assuming ideal conditions and excluding the perturbations analyzed in Task 2, this approach isolates the specific influence of varying living standards to facilitate a definitive baseline comparison.

5.4.1 Demand Scaling Across Comfort Tiers

Minute variations in the comfort factor κ trigger significant shifts in logistical scale. The water demand metrics for the three simulated tiers are summarized in Table 7.

Table 7: Water Demand Metrics under Different Comfort Scenarios

Demand Tier	κ	**Water Inventory (metric tons)**	**Daily Rep. (metric tons)**	**Initial Trans. (metric tons)**	**Annual Total (metric tons)**
Survival Standard	1	290	29.1	1,163	10,622
Comfort Standard	50	2,250	225.1	9,003	82,162
Luxury Standard	250	10,250	1,025.1	41,003	374,162

Table 7 illustrates a multiplier effect: as the standard shifts from Survival to Luxury, the annual requirement surges by over 35-fold. Supported by high-efficiency recycling, the daily replenishment required from Earth remains a small fraction of the total demand, emphasizing the value of closed-loop systems.

5.4.2 Integrated Trade-off Analysis of Transport Schemes

To satisfy these requirements, we mapped the demand indicators onto the transport framework of Model I. Four distinct schemes were evaluated based on their baseline capacities, as defined in Table 8.

Table 8: Definition and Baseline Capacity of Transport Schemes

Scheme	Description	Daily Capacity (t/d)	Specific Energy (GJ/t)
1	Space Elevator Only (3 units)	1,471	157.2
2	Rocket-Only (10 sites)	1,250	506.1
3	Hybrid (Elevator + 10 sites)	2,721	317.5
4	Hybrid (Elevator + Low-lat.)	1,971	243.5

By consolidating the simulation outputs for all scenarios, Table 9 presents a comprehensive comparison of scheme performance across κ tiers.

Table 9: Consolidated Performance Metrics of Transport Schemes across Comfort Tiers

Scenario	Scheme	Initial Days	Initial En. (TJ)	Monthly Days	Annual En. (PJ)
Survival ($\kappa = 1$)	Scheme 1	0.79	182.8	0.59	1.67
	Scheme 2	0.93	588.6	0.69	5.38
	Scheme 3	0.43	369.2	0.32	3.37
	Scheme 4	0.59	283.1	0.44	2.59
Comfort ($\kappa = 50$)	Scheme 1	6.12	1,415.3	4.59	12.92
	Scheme 2	7.20	4,556.3	5.40	41.58
	Scheme 3	3.31	2,858.1	2.48	26.08
	Scheme 4	4.57	2,191.9	3.43	20.00
Luxury ($\kappa = 250$)	Scheme 1	27.87	6,445.7	20.90	58.82
	Scheme 2	32.80	20,751.2	24.60	189.36
	Scheme 3	15.07	13,016.9	11.30	118.78
	Scheme 4	20.80	9,982.5	15.60	91.09

While water transport is highly feasible for survival, the Comfort Scenario ($\kappa = 50$) serves as the planning representative where efficiency-timeline trade-offs intensify. In this tier, Scheme 1 defines the efficiency baseline, Scheme 3 maximizes throughput to minimize delivery time, and Scheme 4 offers the optimal operational balance. Under the Luxury Scenario ($\kappa = 250$), energy demand for rocket transport peaks at 189.36 PJ. Crucially, even this extreme demand represents only 0.37 percent of the total construction energy, confirming that post-2050 logistical constraints stem from capacity allocation rather than aggregate energy availability.

5.5 Sensitivity Analysis for Water Supply Model

To assess the robustness of Model II and identify the dominant levers in water logistics, we conduct a sensitivity analysis over five variables: recycling efficiency (η), comfort factor (κ), population (N), medical demand parameters, and buffer duration.

1. Parameter Sensitivity Ranking via Tornado Analysis

We systematically varied each parameter within its feasible range while holding others at baseline values ($\kappa = 50$, $\eta = 90\%$, $N = 100,000$, sickness rate $p = 2\%$, buffer $= 30$ days). The resulting impacts on annual water supply are visualized through a tornado diagram (Figure 12).

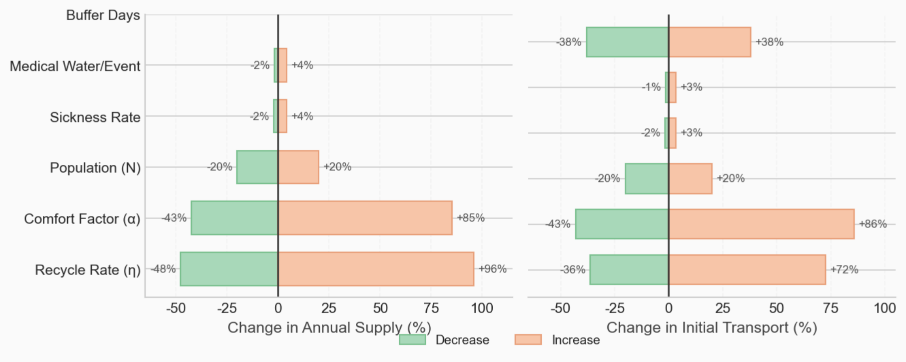

Figure 12: Tornado Diagram

1. **Recycling efficiency** (η) dominates: modest degradations cause large increases in annual resupply because losses scale with $(1 - \eta)$.
2. **Secondary drivers**: comfort (κ) is the next lever (tightening provides headroom), population (N) is near-linear, while medical/buffer settings have limited effect on annual totals but affect initial fill and resilience.

2. Quantitative Sensitivity Coefficients

Table 10 presents the normalized sensitivity coefficients, defined as the percentage change in output per percentage change in input parameter:

Table 10: Sensitivity Coefficients for Key Model Parameters

Parameter	**Baseline**	**Range**	**Δ Annual Supply**	**Sensitivity**	**Rank**
Recycling Rate (η)	90%	80%–95%	-52% to $+96\%$	-9.6	1
Comfort Factor (κ)	50	25–100	-71% to $+17\%$	+0.88	2
Population (N)	100k	80k–120k	$\pm 20\%$	+1.0	3
Sickness Rate (p_{ill})	2%	1%–4%	-0.3% to $+0.6\%$	+0.30	4
Buffer Days	30	15–45	$\pm 0\%$*	0	5

The sensitivity coefficient for η reaches -9.6: a 1% decrease in recycling efficiency increases annual resupply by about 9.6%, making recycling integrity the primary control variable.

3. Worst-Case Stress Testing

Beyond univariate analysis, we evaluate system performance under compound adverse conditions through Monte Carlo-informed scenario construction. Table 11 summarizes nine representative scenarios spanning from optimal to catastrophic conditions:

Table 11: Worst-Case Scenario Analysis: Water Supply System Stress Test

Scenario	κ	η	N **(k)**	p_{ill}	**Buffer**	**Annual (kt)**	**Capacity (%)**	**Feasible**
Survival Mode	1	90%	100	2%	30	14.2	2.7%	✓
Comfort Mode (Baseline)	50	90%	100	2%	30	85.8	16.0%	✓
Luxury Mode	250	90%	100	2%	30	377.8	70.3%	✓
Recycle Degradation	50	80%	100	2%	30	167.9	31.3%	✓
Epidemic Outbreak	50	90%	100	5%	45	100.4	18.7%	✓
Population Surge	50	90%	120	2%	30	102.9	19.2%	✓
Moderate Stress	100	85%	110	3%	40	264.4	49.2%	✓
Worst Case	**250**	**80%**	**120**	**5%**	**45**	**919.8**	**171.3%**	✗

The stress test shows that feasibility is controlled mainly by the joint condition $\eta < 82\%$ **with** $\kappa > 200$, under which annual water demand can exceed elevator capacity (worst case: **171.3%**), whereas moderate stress remains well within capacity (49.2%).

5.6 Conclusion and Strategic Insights

1. **Trade-off between Comfort and Capacity**: Quantitative analysis confirms that quality of life acts as a logistical amplifier. At the Luxury tier, the annual water replenishment occupies 69.68% of the theoretical annual capacity of the space elevator system. Maintaining high living standards significantly constricts the transport window for other critical infrastructure and scientific materials.

2. **Tiered Strategy Recommendations**:

    - **Initial Operations**: We recommend Survival tier combined with Scheme 1. This minimizes energy consumption while ensuring the survival threshold for 100,000 residents, reserving hybrid capacity for urgent infrastructure tasks.
    - **Mature Operations**: We recommend Comfort tier combined with Scheme 4. This provides a balance between speed and cost, completing the initial filling in 4.57 days with a reasonable energy footprint.

3. **System Resilience**: Thanks to the 90% efficient recycling system and the 30-day emergency buffer, the water logistics chain exhibits high resilience. Even during a full month of transport disruption, the survival of the colony remains uncompromised, allowing for logistical maintenance and error recovery.

6 Model Extension

Following the assessment of construction timelines and life-support logistics, this chapter extends the model to incorporate the environmental footprint as a critical optimization boundary. By quantifying the ecological externalities of various transport scenarios, we refine previous strategic selections to promote symbiosis between Earth's conservation and lunar sustainability.

6.1 Quantification of Baseline Environmental Impacts

We extended the environmental sub-module based on the energy conversion logic of Model I. For liquid oxygen/methane (LOX/CH_4) propulsion systems, fuel consumption is mapped to emission mass via the combustion equation:

$$CH_4 + 2O_2 \longrightarrow CO_2 + 2H_2O \tag{11}$$

A life-cycle assessment (LCA), including fuel production and power generation, was conducted for four baseline scenarios, as summarized in Table 12.

Table 12: Environmental Metrics Across Four Baseline Scenarios

Scenario	**Timeline (a)**	**Total Fuel (Mt)**	**Total CO_2 (Mt)**	**Stratospheric H_2O (Mt)**	**Intensity (Mt/yr)**
Rocket-Only	219	9,596	12,968	1,919	59.2
Elevator-Only	186	69	846	0	4.5
Hybrid (Time-Prioritized)	101	4,445	6,414	882	63.7
Hybrid (Balanced)	139	2,438	3,856	477	27.7

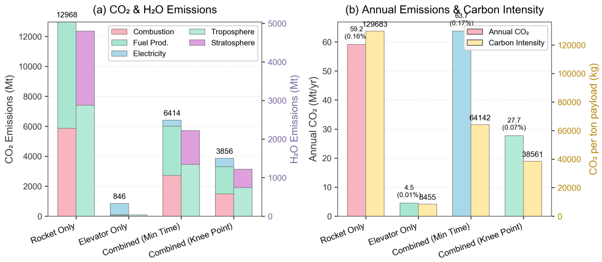

Figure 13: Multi-dimensional environmental impact analysis

The analysis in Figure 13 clearly distinguishes the ecological costs:

- The **Rocket-Only scenario** incurs the highest CO_2 emissions and stratospheric H_2O injection, posing the greatest risk to global climate and the ozone layer.
- The **Elevator-Only scenario** completely circumvents stratospheric pollution, with total CO_2 emissions being only 6.5% of the rocket scenario.
- The **Balanced Hybrid scenario** achieves a 40% reduction in CO_2 and a 46% reduction in stratospheric H_2O compared to the Time-Prioritized plan, successfully equilibrating construction efficiency with environmental load.

6.2 Earth Port Assessment

While the space elevator facilitates zero-emission transit, its Earth Port—a permanent maritime structure in equatorial waters—may impact local marine ecosystems. Drawing on research from the International Space Elevator Consortium (ISEC) [2], we analogized the Earth Port to offshore oil rigs, the most mature existing maritime facilities.

Table 13: Comparative Impact: Earth Port vs. Offshore Oil Platforms

Impact Type	**Oil Platform**	**Earth Port**	**Range**
Physical Damage	Benthic disruption (anchoring)	Anchor zone-only	Localized
Chemical Discharge	Drilling fluids/water	No discharge	Negligible
Interference	Light/acoustic disturbance	Navigation light + low-freq noise	Localized
Habitat Alteration	Artificial reef effect	Hard substrate (sessile colonies)	Positive
Accident Risk	Hydrocarbon spills	No leakage risk	Negligible

The analogy confirms that the Earth Port's environmental footprint is significantly lower than that of existing oil platforms. Compared to the widespread atmospheric damage caused by rockets, the localized impacts of the space elevator are nearly negligible, reinforcing its ecological superiority.

6.3 Environmental Priority Optimization Framework

To prioritize Earth's ecological integrity, we redefine the optimization logic by transitioning from a multi-objective trade-off to an Environmental Single-Objective framework. In this model, environmental impact is the primary objective to be minimized, while construction time and energy consumption are treated as boundary constraints:

$$\min \quad J_{env} = \gamma \cdot \text{Emission}_{\text{CO}_2} + \delta \cdot \text{Emission}_{\text{H}_2\text{O, strat}} \tag{12}$$

Simulation analysis reveals that by setting J_{env} as the sole optimization target, the system naturally eliminates chemical propulsion dependencies. The 186-year mark represents the absolute floor of ecological externalities—the point where stratospheric pollution is zeroed and carbon emissions reach their global minimum. Beyond this threshold, further extending the timeline yields no additional environmental gain, confirming the Elevator-Only scenario as the most sustainable strategy.

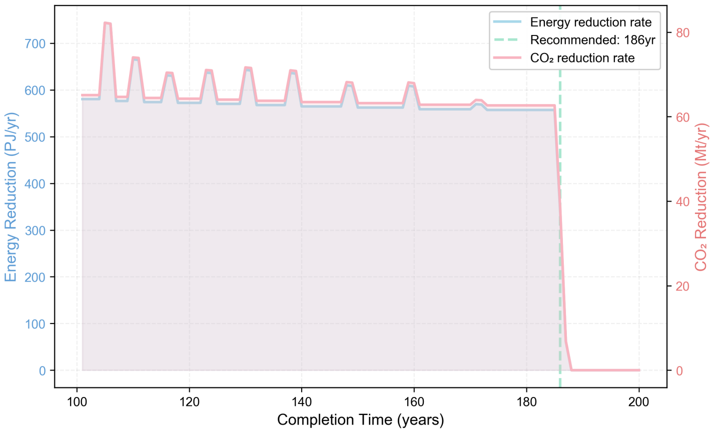

Figure 14: Energy and CO_2 Reduction Potential Analysis.

6.4 Result of Task 4

Conclusively, the 186-year standalone space elevator scenario is proposed as the optimal strategy for lunar colonization. This decision integrates construction duration, energy efficiency, environmental impact, and system reliability.

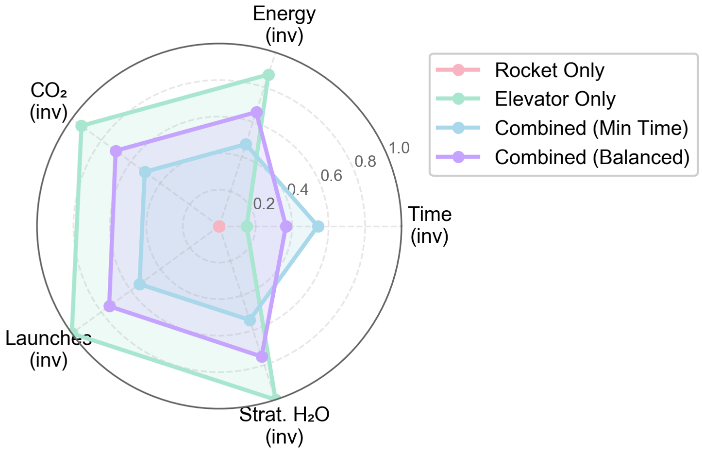

Figure 15: Multi-dimensional Environmental Impact (Radar Chart)

- The rocket-only scenario imposes the most severe environmental risks, with carbon emissions and stratospheric water vapor injection threatening the global climate and ozone stability.
- Furthermore, although the hybrid scenario offers temporal advantages, it incurs substantial emissions and introduces heightened system complexity and failure risks.
- Consequently, the 186-year elevator-only plan represents the most sustainable choice, as it completely eliminates rocket-based pollution by trading temporal efficiency for ecological preservation.

7 Strengths and Weaknesses

7.1 Strengths

- **Comparable cost proxy across technologies.** Using Universal Energy-Equivalent (UEE, Joules) avoids long-horizon monetary uncertainty and enables a thermodynamically consistent comparison between chemical rockets and space-elevator lifting.
- **Interpretable multi-objective decision logic.** The time-opportunity parameter λ converts the energy–time trade-off into a single objective with a clear optimality condition, making stakeholder preferences explicit and auditable.

7.2 Weaknesses and Possible Improvement

- **Parameter uncertainty and long-horizon extrapolation.** Key inputs are necessarily assumed or extrapolated.
- **UEE is not a complete economic cost.** Energy captures a physical lower bound but does not fully represent CAPEX, labor, insurance, and supply-chain constraints.

References

[1] David R Francisco. "NASA-STD-3001 Volume 2, Revision E, NASA Spaceflight Human-System Standard Volume 2: Human Factors, Habitability, and Environmental Health". In: *N/A* A (2025).

[2] Michel Gassent, Kevin Barry, and Peter Swan. *Space Elevators: The Green Road to Space*. ISEC Study Report. Accessed: 2023-10-27. Santa Barbara, CA: International Space Elevator Consortium (ISEC), 2021. URL: `https://isec.org/wp-content/uploads/2021/11/ISEC-2021-Study-Report-Green-Road-to-Space-Final.pdf`.

[3] Gunter's Space Page. *Falcon-9*. `https://space.skyrocket.de/doc_lau/falcon-9.htm`. Accessed: 2026-02-01. Gunter's Space Page, 2026.

[4] Hong Liu et al. "Review of research into bioregenerative life support system(s) which can support humans living in space". In: *Life Sciences in Space Research* 31 (2021), pp. 113–120. ISSN: 2214-5524. DOI: `https://doi.org/10.1016/j.lssr.2021.09.003`. URL: `https://www.sciencedirect.com/science/article/pii/S2214552421000663`.

[5] OpenAI. *GPT*. 5.2. 2026-02-01. 2026.

[6] Zephyr Penoyre and Emily Sandford. *The Spaceline: a practical space elevator alternative achievable with current technology*. 2019. arXiv: `1908.09339 [astro-ph.IM]`.

[7] Peter A Swan et al. "Space Elevators: An Assessment of the Technological Feasibility and the Way Forward". In: *International Academy of Astronautics* (2013).

[8] Mihriban Whitmore, Jennifer Boyer, and Keith Holubec. "NASA-STD-3001, space flight human-system standard and the human integration design handbook". In: *Industrial and systems engineering research conference*. JSC-CN-25695. 2012.

[9] xan3011. *Airline Data Project (MIT 1995-2019)*. `https://www.kaggle.com/datasets/xan3011/airline-data-project-mit-1995-2019`. Accessed: 2026-02-03. 2019.

To: Moon Colony Management (MCM) Agency
From: Team 2618656
Subject: Strategic Recommendation for the Moon Colony Transport and Logistics
Date: February 4, 2026

Dear Members of the MCM Agency Board,

We submit this recommendation for the 100,000-person Moon Colony, developed via rigorous modeling and risk analysis. Below is our actionable strategy, supported by key metrics and an implementation roadmap.

1. Recommended Decision

Our **Phased Hybrid Strategy** utilizes the Space Elevator for primary throughput and low-latitude rockets for resilience. We recommend a **70–80% elevator and 20–30% rocket** split (modeled at 74.6%/25.4%), utilizing dynamic rebalancing to ensure system stability..

2. Core Performance Metrics

The hybrid strategy is Pareto-superior, delivering 100Mt within 155–160 years (95% confidence). It requires 38,750 PJ—a 23.4% energy saving—and reduces cumulative CO_2 by ~40% compared to rocket-only alternatives, significantly mitigating environmental risks.

3. Risk Controls & Triggers

Resilience is embedded through the following operational protocols:

- **Capacity:** ≥10% reserve at all Harbours; <14-day recovery target.
- **Surge Policy:** Shift to 30–35% rocket share if 180-day elevator output <90% nominal.
- **Logistics:** Prioritize low-latitude launches for emission efficiency and redundancy.

4. Sustainment Logistics (Water Requirements)

Post-habitation annual water needs: 10.6 kt (survival) to 374 kt (luxury). Capacity allocation is critical:

- Mandate ≥85% water recycling efficiency (below this is mission-risk, requiring immediate action).
- Launch with survival-tier standards; elevate comfort only as ISRU and recycling stabilize.
- If low recycling efficiency and high demand coincide, roll back comfort and surge maintenance.

5. Implementation Roadmap

The Moon Colony is feasible within 155–160 years, with significant energy and environmental savings vs. rocket-only. Success depends on the hybrid architecture, paired with resilience mandates and auditable triggers to control schedule risk.

We stand ready to provide additional support.

Respectfully submitted,
Team 2618656

Report on Use of AI

1. OpenAI ChatGPT (accessed 2026–02–01)

Purpose: We used a generative AI assistant only for (i) outlining and language polishing, (ii) checking for clarity/consistency of model descriptions, and (iii) LaTeX presentation suggestions (tables/figures). We did **not** use AI to generate final numerical results, figures, or citations; all computations and plots were produced by the team using our own scripts and datasets, and all claims were verified by the team before inclusion.

Where AI was potentially useful in the paper: background framing (Introduction), simplifying and explaining assumptions (Assumptions), verifying the expression/units of key equations (Model I & II), drafting/condensing the robustness narrative (stochastic/CVaR), and improving readability of the recommendation/strengths-weaknesses sections.

Query 1: "Rewrite the following paragraph in concise academic English, keep the meaning, do not add new claims: *Monetary costs are ill-conditioned for an Earth–Moon logistics program spanning many decades due to inflation, technology shocks, and geopolitical uncertainty. We therefore use energy consumption as a stable proxy for marginal cost, enabling a consistent comparison between rockets and the Space Elevator System.*"

Output 1: "Because monetary costs over multi-decade programs are highly uncertain (inflation, technology shocks, and geopolitics), we adopt energy consumption as a physics-grounded proxy for marginal cost, enabling consistent comparison between rockets and the Space Elevator System."

Query 2: "Give a short, publication-style definition of CVaR at confidence level α and one sentence on how it is used in schedule-risk planning."

Output 2: "Conditional Value-at-Risk (CVaR_α) is the expected value of a random variable conditional on outcomes in the worst $1-\alpha$ tail of its distribution. In schedule-risk planning, CVaR_α summarizes "expected worst-case" delays beyond the α-quantile and can be penalized to obtain risk-averse plans."

Query 3: "Provide a minimal LaTeX example (xcolor + booktabs) to highlight one recommended row in light blue and keep the rest plain, suitable for an MCM paper table."

Output 3: "Use `\usepackage[table]{xcolor}` and define a color, then apply `\rowcolor{...}` to the recommended row. Example:

```
\definecolor{recommendColor}{RGB}{217,234,249}
\begin{tabular}{lccc}
\toprule
... \\
\midrule
Baseline & ... \\
\rowcolor{recommendColor} Recommended & ... \\
\bottomrule
\end{tabular}
```

2. Cursor IDE Autocomplete (Free plan, accessed 2026–02–01)

   Purpose: Auto-completions for code used in preparing our models.