Dynamic Reserves, Surge Staffing, and Flexible Personnel Allocation in Complex Spaceflight Missions

Zachary Barnes

March 13, 2026

Executive Summary

This report evaluates the development of a first-of-a-kind, one-of-a-kind spaceflight mission using a dynamic spreadsheet model that integrates uncertainty, schedule risk, and management flexibility. The project is represented as a dual-track technology development effort with separate Technology Readiness Level (TRL) trajectories for a cruise-stage subsystem and a lander-stage subsystem. Launch occurs only after both tracks reach the required pre-launch maturity, after which the mission enters a post-launch cruise phase and then an operations phase. The model evaluates mission value using discounted cash flow net present value (NPV), allowing the effects of timing, development cost, delay, and flexibility to be compared on a consistent basis.

The model includes uncertainty in annual budget authority, technical progress, execution cost growth, allocation instability between the two technology tracks, and discrete TRL transition delays - all modeled in a Monte Carlo simulation. These uncertainties create both cost dispersion and schedule dispersion. The resulting mission outcomes are therefore highly path-dependent: in some futures, the mission converges efficiently, while in others, a lagging subsystem, poor budget year, or repeated delay event forces a long and expensive tail.

Several flexibilities were implemented directly in the model. These include bounded TRL-based annual reallocation of resources between the cruise-stage and lander-stage workstreams, and a reserve mechanism that allows unspent budget authority to be carried forward and used in later shortfall years. A surge-staffing flexibility has also been added, meaning when realized subsystem lag exceeds a threshold, staffing can temporarily rise above the baseline plan, subject to an affordability constraint, a staffing cap, and a cap on how much extra staffing can accelerate TRL progress in a single year. These flexibilities are not merely spreadsheet devices; they map directly to real management and technical decisions such as reallocating engineers and test resources, maintaining management reserve, and preserving optionality in technical execution.

This model captures the fact that management can sometimes convert favorable funding conditions into future schedule protection and can selectively apply additional staffing when technical lag becomes meaningful.

The principal recommendation of this report is that a mission of this type should not be managed as a rigid, single-path development program. Instead, the program should be planned with explicit flexibility in both technical execution and time-based financial management. In particular, management should preserve a dynamic reserve, maintain formal reallocation and surge rules tied to subsystem maturity gaps, and monitor the two TRL tracks independently rather than collapsing them into a single maturity metric. A disciplined baseline plan remains necessary, but the expected value of the mission improves when management has the ability to respond to uncertainty as it is revealed.

Contents

1 Introduction

The purpose of this analysis is to demonstrate how flexibility and real options can be used to improve the management of a first-of-a-kind, one-of-a-kind spaceflight mission under uncertain conditions. Whereas recurring projects may be forecast using prior analogs, a novel space mission often combines bespoke subsystem development, uncertain budget conditions, and incomplete technical maturity at program start. These features make both cost and schedule difficult to forecast and make rigid planning especially vulnerable to downside outcomes.

1.1 Motivation

The motivation for this project is twofold. First, large and technically ambitious spaceflight programs frequently experience cost growth and schedule delay because critical technologies do not mature uniformly across the mission architecture. Second, management decisions made in response to those delays are often informal, reactive, or insufficiently tied to actual information revealed during development. This project seeks to create a model that can both capture those uncertainties and evaluate which flexibilities are most useful in practice.

A further motivation is methodological. Many real options studies focus on infrastructure or long-life systems where flexibility is represented primarily as a capacity expansion or timing choice. [1] In contrast, this project focuses on a development program itself. The “system” being managed is the mission development process: staffing, subsystem maturity, annual budget authority, reserve, launch readiness, cruise, and operations. The objective is therefore not only to forecast mission cost and timing, but to identify management actions that improve expected mission value.

1.2 Scenario

The scenario considered here is a notional first-of-a-kind, one-of-a-kind spaceflight mission with two major subsystem development tracks:

- **Cruise-Stage Subsystem:** which enables travel between destinations
- **Lander-Stage Subsystem:** which is primarily responsible for the entry, descent, and landing operations once arrived at the mission destination.

This configuration is standard for many interplanetary missions, such as the Spirit & Opportunity Mars Exploration Rovers (MER) - illustrated in Figure 1, and InSight - a NASA Jet Propulsion Laboratory spacecraft sent to the Martian surface in 2018. [4]

The modeled mission targets a nominal 10-year completion but is presented on a 15-year time horizon so that uncertainty-driven schedule slip can be observed. After launch, the mission is expected to enter a post-launch cruise period, followed by an operations period.

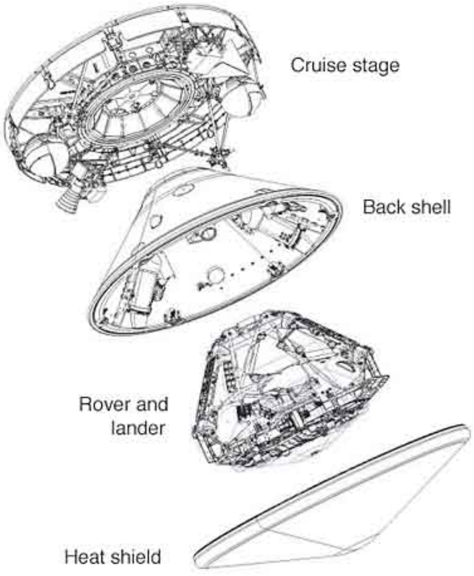

Figure 1: MER Spacecraft Subsystem [4]

1.3 System Design

The primary driver of performance in this system is Technology Readiness Level (TRL), which is used as a proxy for overall subsystem maturity throughout the mission lifecycle. [2]

Each track begins the project at its own initial TRL and progresses annually according to planned targets, realized staffing, uncertainty in development progress, and the potential occurrence of delay events. Because the mission requires both elements to be sufficiently mature before launch, the overall schedule is paced by the slower of the two tracks.

Launch readiness is triggered once both technology tracks cross the required threshold (in this case, TRL 8). The mission then enters a fixed-duration 2-year cruise phase

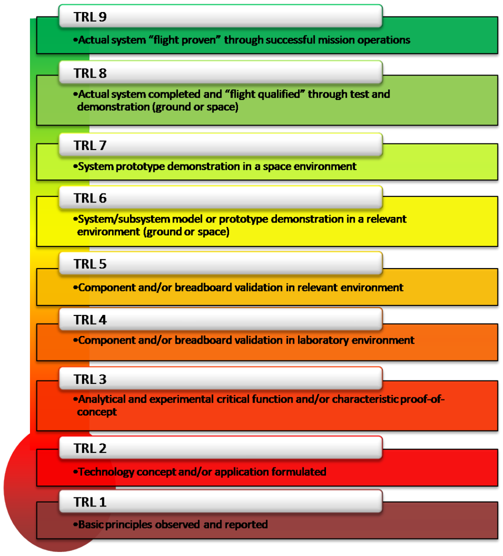

Figure 2: NASA Technology Readiness Level (TRL) Definitions [2]

followed by a one-year operations phase. Development costs continue to accrue until technical completion and associated shutdown conditions are met.

The deterministic model is organized in annual fiscal-year increments. For each year, the model:

1. determines TRL progression through a deterministic demand model
2. matches staffing levels to a deterministic demand curve-driven TRL progression
3. allocates staffing between cruise-stage and lander-stage development based on a fixed ratio
4. computes annual labor, non-labor, & overhead costs
5. calculates annual cash flow and discounted present value.

1.4 System Utility

The mission's utility is represented as discounted net present value, where the NPV reflects the total cost of the mission. The mission assumes a real discount rate of 1.9%, similar to the rate typically allocated to federally-funded large-scale government projects [5]. The key economic intuition is straightforward: value increases (cost decreases) when mission benefits are realized sooner and when development costs are incurred later or avoided altogether. The model, therefore, values both schedule acceleration and cost avoidance. Any flexibility that helps the mission reach launch readiness sooner, preserve staffing continuity during bad budget years, reduce the cost of a long program tail, or avoid unnecessary post-completion spending will tend to improve NPV.

1.5 System Constraints

Several constraints shape the development problem:

- launch cannot occur until *both* cruise-stage and lander-stage maturity are sufficient;
- annual staffing is capped, and follows a ramp-up/ramp-down based on a deterministic demand profile
- budget authority is exogenous and uncertain
- resource reallocation between workstreams is bounded
- operations cannot begin until after launch and post-launch cruise phases are complete

These constraints are what make flexibility valuable: management cannot simply assume a deterministic TRL progression, unlimited access to staffing, instantaneous reallocation, or an on-schedule launch.

2 Baseline Analysis

The baseline case represents the mission under planned staffing, planned allocation, and nominal development assumptions. The mission is evaluated across a 10-year horizon.

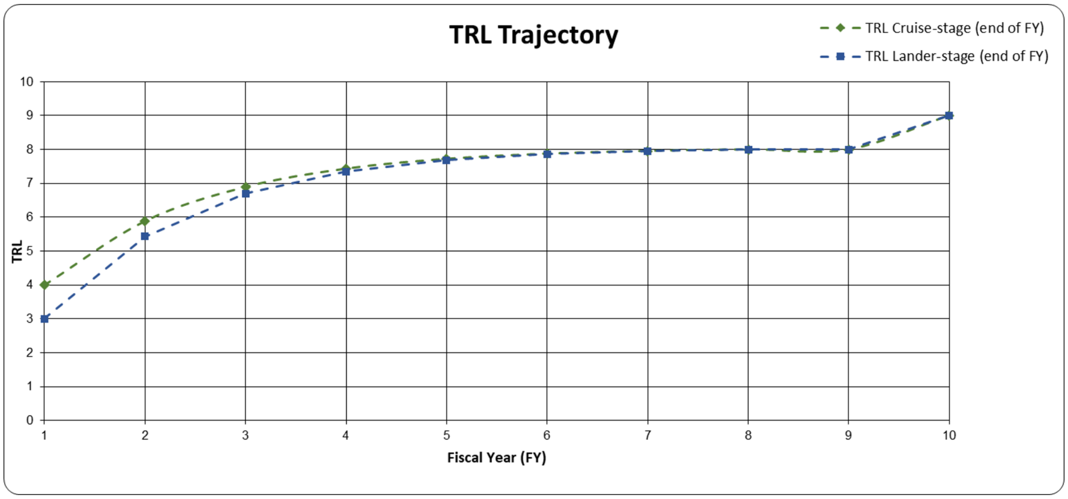

Figure 3: Subsystem TRL Progression (Deterministic Case)

In this initial instance of the model, the TRL progression is deterministic and is forecast using a dynamic progression model given in the equation below:

$$TRL(t) = TRL_{final} - \alpha e^{-\beta t} \tag{1}$$

where $TRL(t)$ = the TRL in the fiscal year t

α = Initial TRL ($\alpha_{cruise} = 3$; $\alpha_{lander} = 4$)

β = growth/convergence rate parameter ($\beta_{cruise} = 0.626$; $\beta_{lander} = 0.585$)

Each subsystem TRL will remain at TRL 8 until the mission operations phase commences after the two year journey, at which point TRL 9 will be achieved, and the mission concludes.

Table 1: Baseline Analysis Results

Metric	Value
Baseline NPV	-\$1,053M
Nominal/Expected launch year	FY 8
Mission completion year	FY 10 (EOY)

This case remains useful because it provides a reference point against which both uncertainty and flexibility can be judged. Without that reference, it is difficult to distinguish whether a flexibility genuinely improves mission value or simply shifts cost and timing around the model.

3 Sensitivity

A one-way sensitivity analysis is used to identify which uncertainties and flexibility settings have the greatest effect on mission NPV. This step is important because not all uncertain variables deserve equal management attention. Some variables primarily affect cost, others primarily affect schedule, and some create asymmetrical downside risk.

The tornado analysis should be interpreted in two ways. First, it identifies the variables that most strongly influence mission value. Second, it helps distinguish between variables that management can only monitor and variables that management can actively influence.

Table 2: Sensitivity Analysis Inputs and Results

Parameter	Expected	High	NPV at High	Low	NPV at Low
Loaded labor rate	$0.25M	$0.30M	-$1,172M	$0.20M	-$935M
Annual Overhead	$50M	$60M	-$1,145M	$40M	-$961M
Non-Labor Cost Multiplier	2x	2.4x	-$1,132M	1.6x	-$974M
Discount Rate	1.9%	1.52%	-$1,070M	2.28%	-$1,037M

Note that the baseline static NPV used for comparison in this sensitivity analysis is: **-$1,053M**

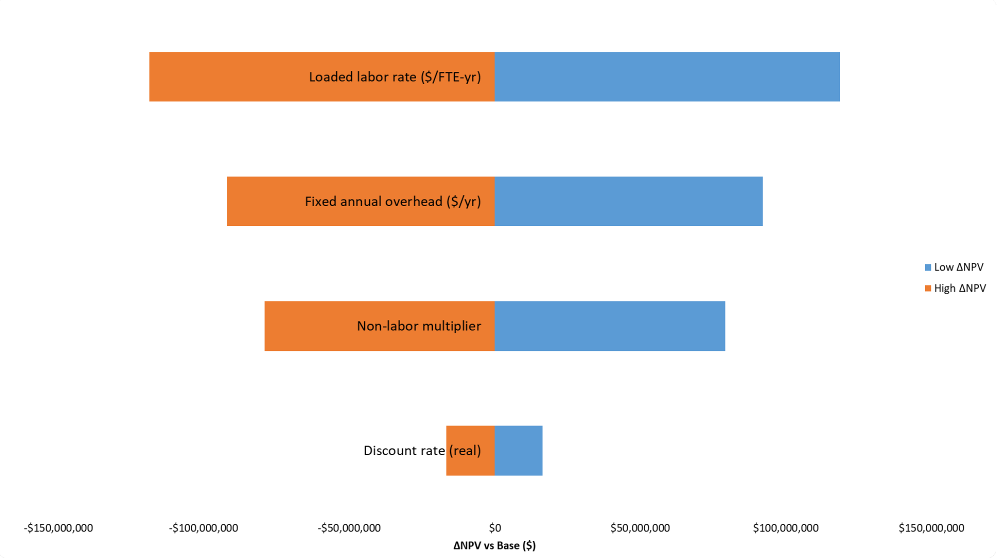

Figure 4: Sensitivity Analysis Results

In general, the model indicates that economic execution variables such as labor rate, overhead, and non-labor burden remain important, but schedule-related technical uncertainties such as TRL volatility, delay probability, and budget shortfall are especially influential because they can extend the project tail and delay benefit realization.

4 Uncertainties

When re-evaluating the model to reduce reliance on deterministic predictions, the model explicitly incorporates uncertainty in both technical progress and financial conditions.

4.1 Annual Budget Volatility

Annual budget authority is modeled as uncertain from year to year. This represents the reality that program funding may vary due to external appropriation decisions, internal portfolio tradeoffs, or changing agency priorities. Budget volatility directly affects how much staffing can be afforded and, therefore, how much productive development work can be completed in a given year.

4.2 TRL Progress Volatility

Even for a given staffing level, realized technical progress does not follow a deterministic path. The model captures this with year-to-year TRL progress volatility. This reflects the reality of rework, uneven subsystem integration, unexpected technical findings, and execution variability.

4.3 Cost Multiplier Volatility

Execution cost is not fixed. The model includes volatility in the cost multiplier to represent variation in procurement cost, vendor performance, rework burden, overhead efficiency, and execution effectiveness. This uncertainty changes how much technical progress can be purchased with a given level of budget authority.

4.4 Budget Allocation Volatility

Even if total funding is known, the effective allocation of resources between cruise-stage and lander-stage effort is not perfectly stable. The model includes allocation volatility to capture management thrash, shifting near-term priorities, and imperfect discipline in annual work balancing.

4.5 TRL Delay Probability and Delay Magnitude

In addition to continuous technical variability, the model includes discrete delay events that reduce realized TRL progress in a given year. This is intended to mimic the kinds of development setbacks common in first-of-a-kind programs: failed tests, supplier issues, integration anomalies, or qualification shortfalls.

5 Performance Under Uncertainty

The uncertainty case represents the mission under realized staffing, planned allocation, and stochastically-determined development outcomes, subject to the uncertainties defined above. The 15-year horizon allows for visibility of performance in runs that extend past the 10-year nominal schedule.

Figure 5: Mission Lifecycle Progress Example with Uncertainty

To understand how the uncertainties interact in combination, the mission is evaluated with Monte Carlo simulation. Monte Carlo simulation was used to quantify how the mission behaves when all of the modeled uncertainties act at the same time rather than one at a time. In the spreadsheet, the Randomized NPV Simulation sheet generates 2,000 simulation runs, with each run corresponding to one full realization of the mission development path under uncertainty.

In each run, the model resamples annual budget volatility, Technology Readiness Level (TRL) progress, cost multiplier, & budget allocation volatility, as well as discrete delay events. It then propagates those through staffing, subsystem maturation, launch & operations timing, cost, and finally net present value (NPV). This produced a population of 2,000 possible NPV outcomes rather than a single deterministic estimate. Rather than focusing on a single expected path, the simulation reveals the distribution of plausible NPVs and schedule outcomes.

This is especially important in the present case because several uncertainties interact nonlinearly. A poor budget year can reduce staffing, reduced staffing can lower realized TRL growth, slower TRL growth can delay launch, and delayed launch pushes benefit farther into the future while increasing the amount of discounted cost carried by the program. The result is a distribution with meaningful downside tail risk.

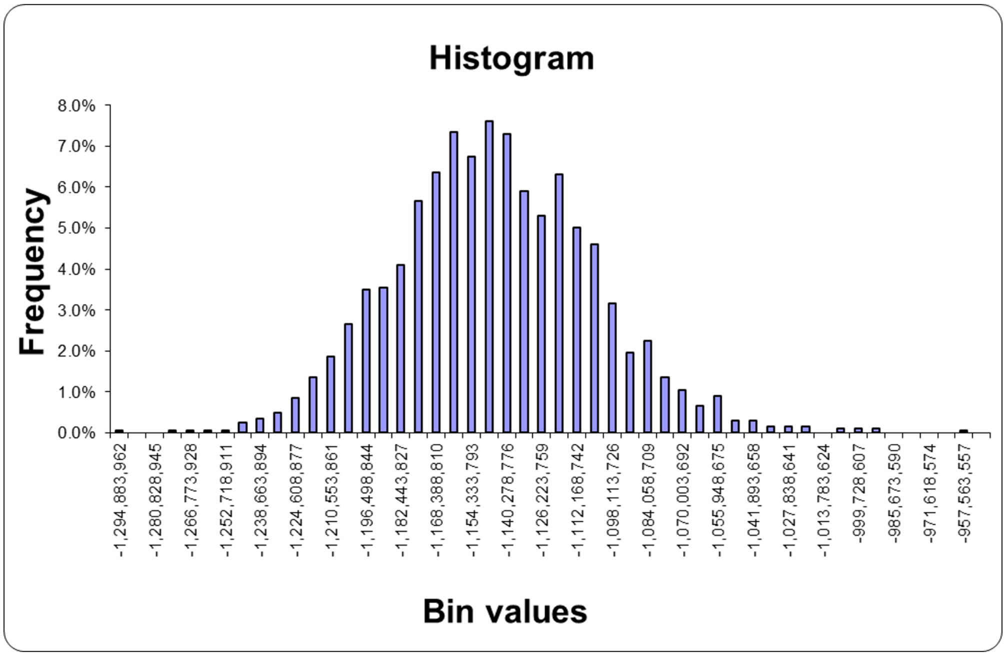

Figure 6: Histogram of Mission NPV Distribution with Uncertainty

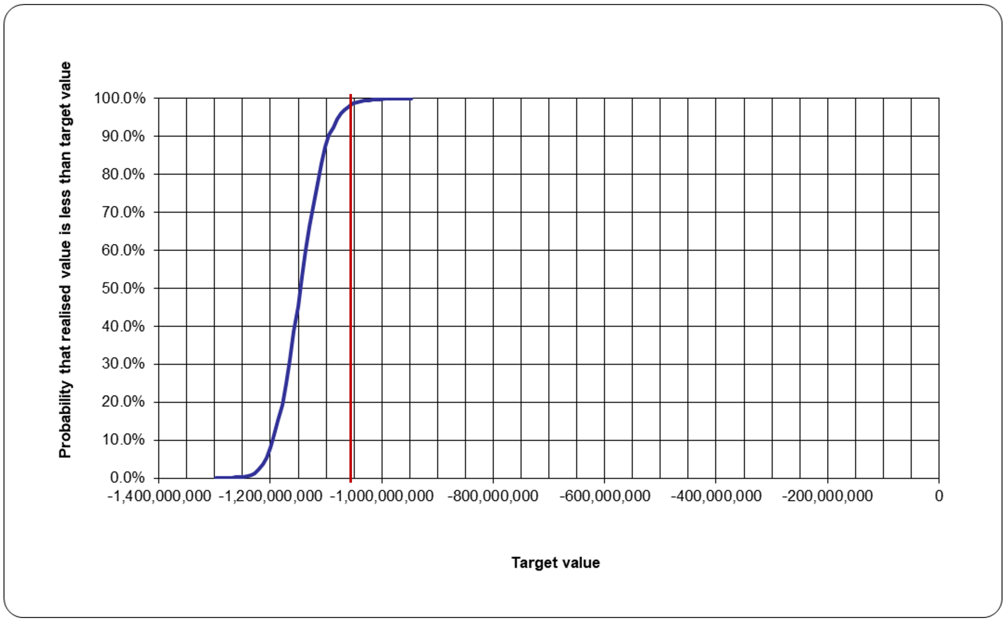

Figure 7: Approximate Cumulative Distribution Function of Mission NPV with Uncertainty

The histogram and cumulative distribution function (CDF) were then used to interpret that simulated outcome set. The histogram on the simulation sheet groups the 2,000 NPV results into 49 intervals (bins), using the simulated minimum and maximum NPV values to define the lower bound and step size of each interval. This makes it possible to see where outcomes cluster most heavily, how wide the spread is, and whether the distribution is skewed toward downside or upside cases.

The CDF uses the same simulated sample set and bin structure to show the cumulative probability of achieving a value less than or equal to a given NPV threshold. That makes it possible to read off probabilities directly, such as the chance of a negative NPV or the approximate P10, P50, and P90 outcomes. A similar approach was used to extract the chance of the mission launching on, or ahead of the FY 8 target.

Table 3: Performance Under Uncertainty

Metric	Value
Mean NPV	-\$1,146M
Standard Deviation	\$41M
P10 / P50 / P90 NPV	-\$1,196M | -\$1,147M | -\$1,094M
Probability of launch by target year	0%

The higher NPV illustrated in this uncertainty analysis does not indicate that the system is performing more poorly than in the deterministic case, but rather that the initial static analysis was insufficient in capturing the true range of factors that affect the true system performance outcome.

Another noteworthy observation is that when accounting for the uncertainty present in the system performance, the probability of a launch by the FY 8 target year drops to **0%**.

The purpose of this section is not solely to state the higher average NPV and inability to launch on target. It is to understand the shape of the uncertainty. The left tail of the distribution corresponds to cases in which budget shocks, cost shocks, and repeated TRL delays interact in a way that drives long development tails and weaker mission value. The right tail corresponds to cases in which both technology tracks converge efficiently and the project enters operations sooner than expected.

6 Real Options

This analysis considers several real options intended to reduce downside and improve expected mission value. The implemented flexibilities fall into two broad categories: technical/programmatic options tied to subsystem maturity, and time-based financial options tied to reserve and schedule management.

6.1 Real Options for Technical Uncertainties

The first major flexibility is bounded annual reallocation between cruise-stage and lander-stage development. If one subsystem lags the other materially, a portion of annual effort can be shifted toward the lagging track. In the spreadsheet model this appears as a maximum allocation shift per year and an allocation rule tied to maturity gaps. In management terms, this corresponds to reallocating cross-trained engineering labor, system integration attention, and technical resources toward the subsystem that is pacing launch readiness.

The second major technical option is surge staffing. If the prior fiscal year reveals that one or both tracks are lagging their target TRL by more than a management-defined threshold, the program can temporarily raise the effective staffing cap above the baseline plan. This does not imply unlimited acceleration. The model includes a separate cap on how much extra staffing can increase the TRL closing rate in a single year, reflecting real bottlenecks such as facility throughput, test cadence, integration bandwidth, and diminishing returns from adding personnel.

The third key technical option is the explicit dual-track TRL structure itself. By managing and evaluating cruise-stage and lander-stage maturity separately, management can identify the pacing subsystem instead of hiding the problem inside an average maturity metric. This improves observability and enables action.

Table 4: Technical Real Options

Option	Description
Bounded reallocation	Shift a limited fraction of annual effort toward the lagging subsystem when maturity gaps exceed a threshold.
Surge staffing	Temporarily raise the effective staffing cap when subsystem lag exceeds a trigger threshold, subject to affordability and a bounded TRL acceleration effect.
Dual-track maturity monitoring	Maintain separate cruise-stage and lander-stage TRL trajectories to preserve visibility into the pacing item.

These are genuine management and technical options because they can be exercised in response to information revealed during the program, rather than being fixed design decisions made once and for all at project start.

6.2 Real Options for Time-Based Parameters

The principal time-based real option in the model is the dynamic reserve. In years where NASA budget authority exceeds productive annual spend, unspent funds are carried into reserve. In later years when authority is low, the reserve can be drawn down to support staffing and preserve progress. This reduces the fragility of the program to single-year funding shocks.

Table 5: Time-Based Real Options

Option	Description
Dynamic reserve	Carry unspent authority forward and deploy it during later shortfall years to preserve continuity.

7 Performance with Flexibilities

Table 6: Performance with Flexibilities

Metric	Value
Mean NPV	-$1,022M
Standard Deviation	$108M
P10 / P50 / P90 NPV	-$1,161M | -$1,023M | -$882M
Probability of launch by target year	58%
Cost of implementation	$625,000
Value of option	$124M

With the aforementioned flexibilities enabled, the NPV improved, on average, by $124M. The probability of the mission launching by the FY 8 target year has also increased to 58%. This accounts for all launch instances occurring at the FY 8 target and also those that occur ahead of schedule, such as the example illustrated in Figure 8.

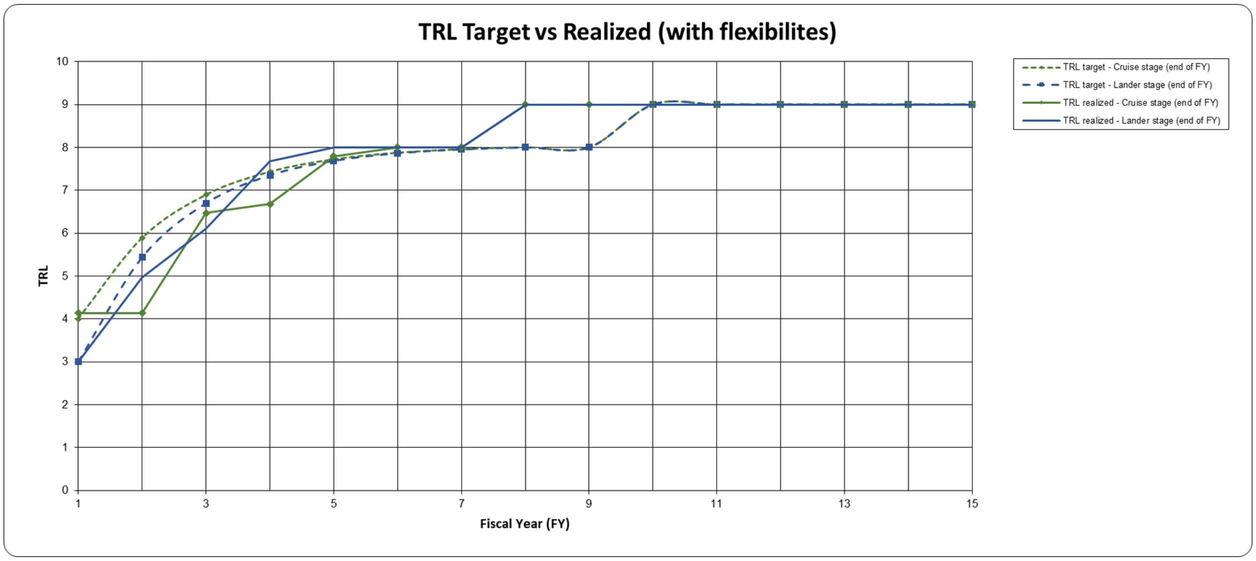

Figure 8: Mission Lifecycle Progress Example with Flexibilities

It should be noted that schedule-oriented flexibilities only create meaningful value when the model allows those flexibilities to affect the pacing mechanism itself. In practical terms, this means that reserves are more valuable when they can preserve staffing continuity, and surge staffing is only valuable when extra staffing can translate into faster TRL closing rather than simply preventing collapse.

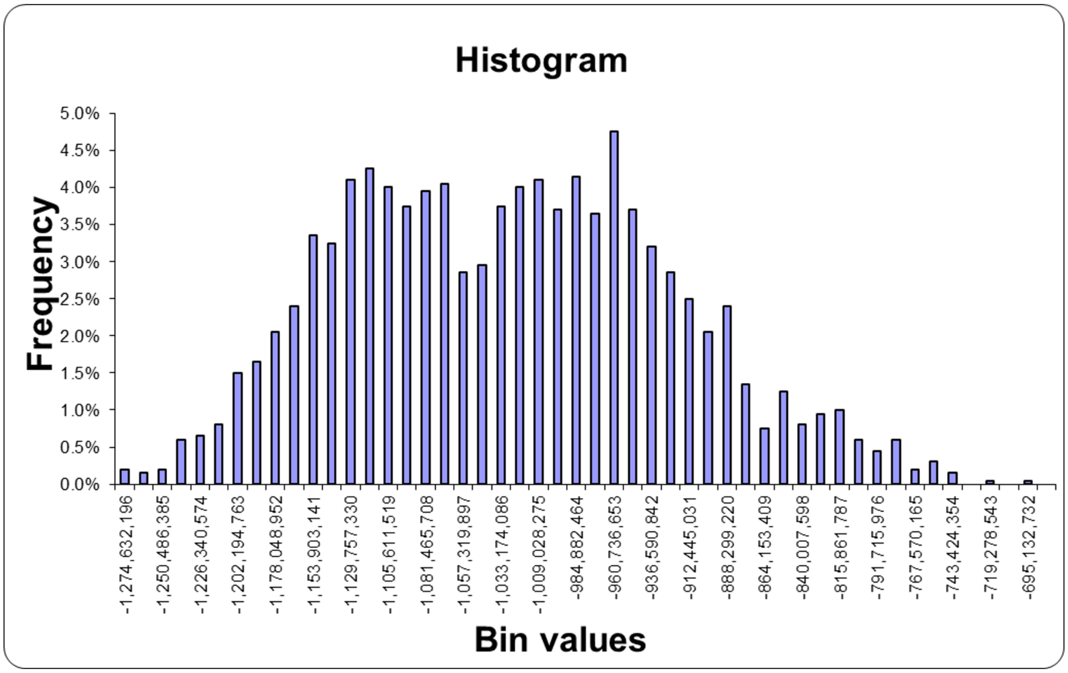

Figure 9: Approximate Cumulative Distribution Function of Mission NPV (with flexibilities)

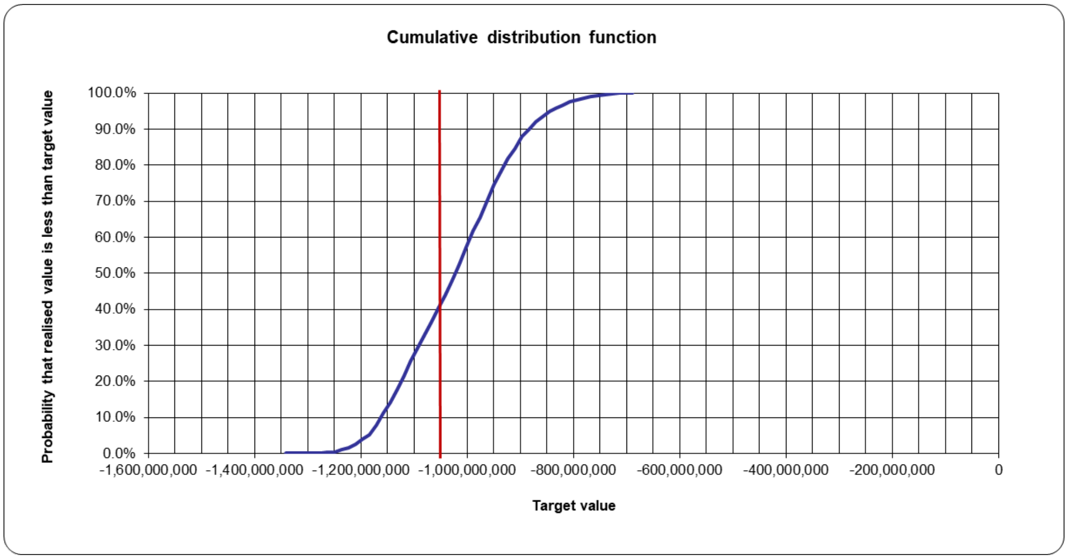

Figure 10: Approximate Cumulative Distribution Function of Mission NPV (with flexibilities)

This is the practical value of the model: it transforms uncertainty from a descriptive risk statement into a set of conditional rules that management can use to act.

8 Lessons Learned & Recommendations

Based on the modeled uncertainties, the flexibilities implemented, and the simulation results, several clear lessons emerged from this project. Most importantly, a first-of-a-kind, one-of-a-kind spaceflight mission cannot be managed effectively with a single deterministic cost and schedule plan. Mission outcomes are driven by the interaction of multiple uncertainties acting at the same time, including annual budget volatility, Technology Readiness Level (TRL) progress volatility, cost growth, allocation instability, and discrete delay events. The analysis therefore suggests that the most effective management approach is neither purely rigid nor unconstrainedly flexible, but rather a disciplined baseline plan supported by explicit option triggers and practical mechanisms for exercising flexibility.

These lessons translate into the following recommendations:

1. **Adopt a dual-track management structure.** One of the clearest lessons from the model is that technical maturity should be tracked at the subsystem level rather than through a single aggregate readiness metric. Modeling separate cruise-stage and lander-stage TRL trajectories showed that the mission is paced by the slower of the two development paths. Management should therefore monitor the two tracks independently and conduct formal pacing-item reviews at each fiscal year so that attention remains focused on the true schedule-driving subsystem rather than on average progress across the mission.

2. **Cross-train technical staff.** The model showed that flexibility has value only when it can be exercised efficiently in practice. In particular, reallocation and surge staffing are only useful if personnel can move across subsystem boundaries without excessive ramp-up losses. Cross-training increases the practical value of both reallocation and surge because it reduces any productivity penalty associated with shifting people between workstreams.

3. **Implement universally shared technical assets.** A related lesson is that organizational and infrastructure choices can materially increase the effectiveness of flexibility. Subsystem interfaces should be standardized where practical so that technical assets such as mechanical ground support equipment (MGSE), electrical ground support equipment (EGSE), tooling, and instrumentation can be reused across workstreams. These actions do not reduce uncertainty directly, but they improve the program's ability to respond efficiently once uncertainty is realized.

4. **Institutionalize bounded annual reallocation.** The analysis showed that reallocation can improve outcomes, but only when it is governed by explicit decision rules. Management should define in advance how many resources can be shifted between workstreams in a single year and what maturity gap justifies the shift. This avoids both underreaction and overreaction, and helps ensure that reallocation is treated as a deliberate real option rather than an sudden response.

5. **Maintain a protected reserve.** One of the strongest financial lessons from the model is that reserve has meaningful value when budget authority is uncertain. Unspent funds in strong years should not be automatically consumed simply because they are available. Instead, they should be carried forward and explicitly protected for years in which budget authority would otherwise be insufficient to sustain critical work. In this way, reserve acts as a stabilizing mechanism that reduces the likelihood of unnecessary schedule slip.

6. **Use surge selectively rather than permanently.** The model suggests that surge staffing can be valuable, but only when it is tied to meaningful technical lag and real bottlenecks. Surge should therefore be used as a temporary intervention rather than as a standing mode of execution. In practice, this means triggering surge only when subsystem maturity falls substantially behind plan.

7. **Use sensitivity results to prioritize executive attention.** The uncertainty analysis, including the tornado diagram, histogram, Monte Carlo simulation, and cumulative distribution function, reinforced that not all variables impact performance equally. Attention should therefore be concentrated on the parameters identified as dominant drivers of net present value, cost, and schedule. This helps ensure that management effort is directed toward the uncertainties that significantly shape mission performance rather than those that are merely visible or easy to discuss.

Taken together, these lessons suggest that the preferred management strategy is to preserve a disciplined baseline plan while embedding practical, pre-defined flexibility into both technical execution and financial management. Ultimately, the goal should not be to eliminate uncertainty, but to structure the program so that it can respond intelligently when uncertainty is revealed.

9 Conclusion

This report demonstrates how a first-of-a-kind, one-of-a-kind spaceflight mission can be modeled as a development program under uncertainty, and how that uncertainty can be managed using explicit real options. The value of the mission depends not only on the original plan, but also on whether management retains the ability to respond intelligently when technical maturity diverges across subsystems, budgets fluctuate, or execution costs vary.

The most important conclusion from the analysis of this model is that flexibility must be both *funded* and *capable of changing the pacing mechanism.* Dynamic reserve is useful because it allows favorable funding years to protect future continuity, but reserve alone is not enough if additional resources cannot actually accelerate the lagging workstream. Surge staffing, dual-track maturity visibility, and flexible personnel allocation together create a more realistic and more valuable representation, and higher-performing example of how a complex development program should be managed.

In this sense, the model is both an analytical tool and also a template for how to structure real program management decisions in uncertain development environments.

A Model Equations and Notation

The model, its uncertainties, and its flexibilities were modeled within Microsoft Excel [3], using set parameters, uncertainty distributions, and logic-based boolean triggers to implement flexibilities. Below are screenshots that illustrate the three primary iterations of the model:

A.1 Static NPV (Baseline) TRL Progression Mission Lifecycle Model

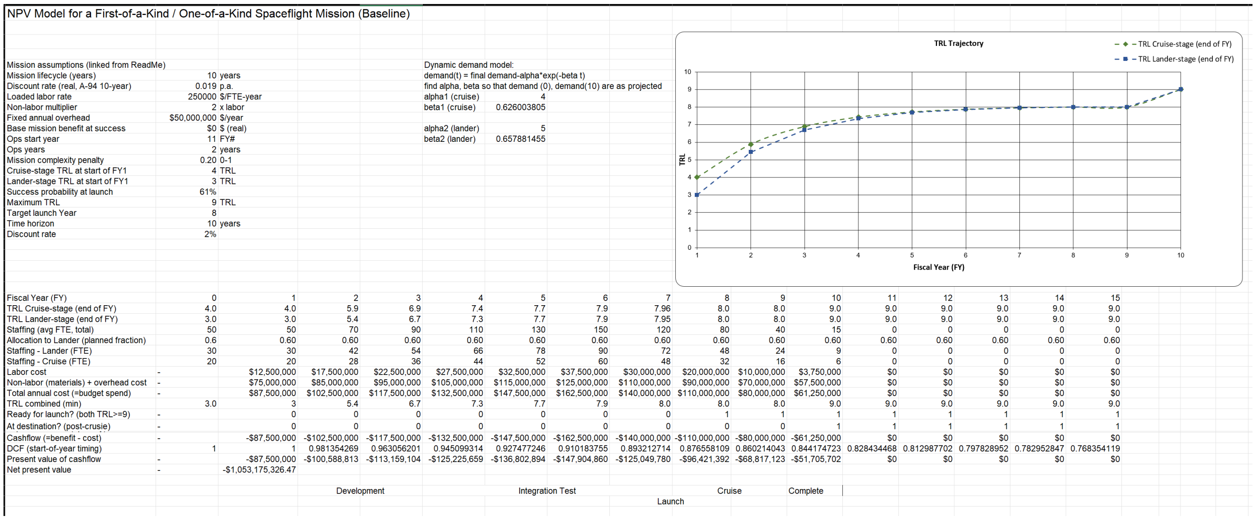

NPV Model for a First-of-a-Kind / One-of-a-Kind Spaceflight Mission (Baseline)

Mission assumptions (linked from ReadMe)		
Mission lifecycle (years)	10	years
Discount rate (real, A-94 10-year)	0.019	p.a.
Loaded labor rate	250000	$/FTE-year
Non-labor multiplier	2	x labor
Fixed annual overhead	$50,000,000	$/year
Base mission benefit at success	$0	$ (real)
Ops start year	11	FY#
Ops years	2	years
Mission complexity penalty	0.20	0-1
Cruise-stage TRL at start of FY1	4	TRL
Lander-stage TRL at start of FY1	3	TRL
Success probability at launch	61%	
Maximum TRL	9	TRL
Target launch Year	8	
Time horizon	10	years
Discount rate	2%	

Dynamic demand model:
demand(t) = final demand-alpha*exp(-beta t)
find alpha, beta so that demand (0), demand(10) are as projected

alpha1 (cruise)	4
beta1 (cruise)	0.626003805
alpha2 (lander)	5
beta2 (lander)	0.657881455

Fiscal Year (FY)	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
TRL Cruise-stage (end of FY)	4.0	4.0	5.9	6.9	7.4	7.7	7.9	7.96	8.0	8.0	9.0	9.0	9.0	9.0	9.0	9.0
TRL Lander-stage (end of FY)	3.0	3.0	5.4	6.7	7.3	7.7	7.9	7.95	8.0	8.0	9.0	9.0	9.0	9.0	9.0	9.0
Staffing (avg FTE, total)	50	50	70	90	110	130	150	120	80	40	15	0	0	0	0	0
Allocation to Lander (planned fraction)	0.6	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60
Staffing - Lander (FTE)	30	30	42	54	66	78	90	72	48	24	9	0	0	0	0	0
Staffing - Cruise (FTE)	20	20	28	36	44	52	60	48	32	16	6	0	0	0	0	0
Labor cost	-	$12,500,000	$17,500,000	$22,500,000	$27,500,000	$32,500,000	$37,500,000	$30,000,000	$20,000,000	$10,000,000	$3,750,000	$0	$0	$0	$0	$0
Non-labor (materials) + overhead cost	-	$75,000,000	$85,000,000	$95,000,000	$105,000,000	$115,000,000	$125,000,000	$110,000,000	$90,000,000	$70,000,000	$57,500,000	$0	$0	$0	$0	$0
Total annual cost (=budget spend)	-	$87,500,000	$102,500,000	$117,500,000	$132,500,000	$147,500,000	$162,500,000	$140,000,000	$110,000,000	$80,000,000	$61,250,000	$0	$0	$0	$0	$0
TRL combined (min)	3.0	3	5.4	6.7	7.3	7.7	7.9	8.0	8.0	8.0	9.0	9.0	9.0	9.0	9.0	9.0
Ready for launch? (both TRL>=9)	-	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
At destination? (post-crusie)	-	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1
Cashflow (=benefit - cost)	-	-$87,500,000	-$102,500,000	-$117,500,000	-$132,500,000	-$147,500,000	-$162,500,000	-$140,000,000	-$110,000,000	-$80,000,000	-$61,250,000	$0	$0	$0	$0	$0
DCF (start-of-year timing)	1	1	0.981354269	0.963056201	0.945099314	0.927477246	0.910183755	0.893212714	0.876558109	0.860214043	0.844174723	0.828434468	0.812987702	0.797828952	0.782952847	0.768354119
Present value of cashflow	-	-$87,500,000	-$100,588,813	-$113,159,104	-$125,225,659	-$136,802,894	-$147,904,860	-$125,049,780	-$96,421,392	-$68,817,123	-$51,705,702	$0	$0	$0	$0	$0
Net present value	-	-$1,053,175,326.47														
			Development			Integration Test			Cruise		Complete					
								Launch								

Figure 11: Static NPV Case

A.2 Uncertainty NPV (Unmitigated) TRL Progression Mission Lifecycle Model

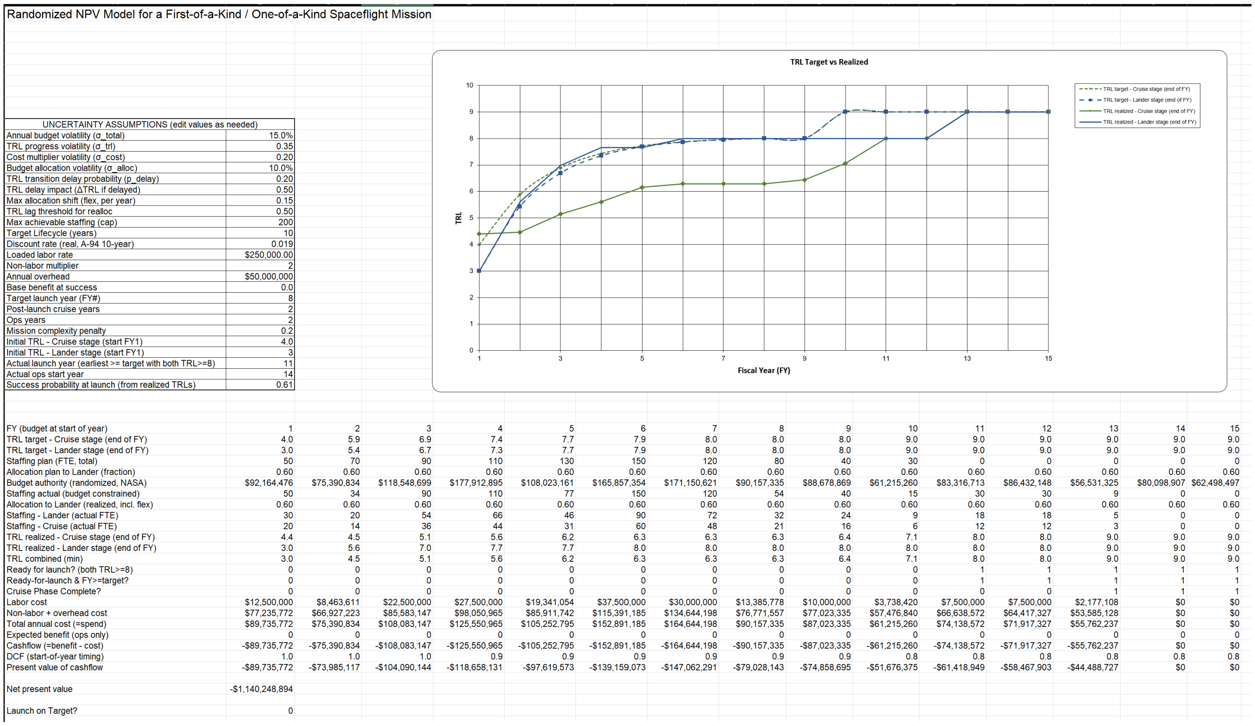

Randomized NPV Model for a First-of-a-Kind / One-of-a-Kind Spaceflight Mission

UNCERTAINTY ASSUMPTIONS (edit values as needed)	
Annual budget volatility (σ_total)	15.0%
TRL progress volatility (σ_trl)	0.35
Cost multiplier volatility (σ_cost)	0.20
Budget allocation volatility (σ_alloc)	10.0%
TRL transition delay probability (p_delay)	0.20
TRL delay impact (ΔTRL if delayed)	0.50
Max allocation shift (flex, per year)	0.15
TRL lag threshold for realloc	0.50
Max achievable staffing (cap)	200
Target Lifecycle (years)	10
Discount rate (real, A-94 10-year)	0.019
Loaded labor rate	$250,000.00
Non-labor multiplier	2
Annual overhead	$50,000,000
Base benefit at success	0.0
Target launch year (FY#)	8
Post-launch cruise years	2
Ops years	2
Mission complexity penalty	0.2
Initial TRL - Cruise stage (start FY1)	4.0
Initial TRL - Lander stage (start FY1)	3
Actual launch year (earliest >= target with both TRL>=8)	11
Actual ops start year	14
Success probability at launch (from realized TRLs)	0.61

FY (budget at start of year)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
TRL target - Cruise stage (end of FY)	4.0	5.9	6.9	7.4	7.7	7.9	8.0	8.0	8.0	9.0	9.0	9.0	9.0	9.0	9.0
TRL target - Lander stage (end of FY)	3.0	5.4	6.7	7.3	7.7	7.9	8.0	8.0	8.0	9.0	9.0	9.0	9.0	9.0	9.0
Staffing plan (FTE, total)	50	70	90	110	130	150	120	80	40	30	0	0	0	0	0
Allocation plan to Lander (fraction)	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60
Budget authority (randomized, NASA)	$92,164,476	$75,390,834	$118,548,699	$177,912,895	$108,023,161	$165,857,354	$171,150,621	$90,157,335	$88,678,869	$61,215,260	$83,316,713	$86,432,148	$56,531,325	$80,098,907	$62,498,497
Staffing actual (budget constrained)	50	34	90	110	77	150	120	54	40	15	30	30	9	0	0
Allocation to Lander (realized, incl. flex)	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60
Staffing - Lander (actual FTE)	30	20	54	66	46	90	72	32	24	9	18	18	5	0	0
Staffing - Cruise (actual FTE)	20	14	36	44	31	60	48	21	16	6	12	12	3	0	0
TRL realized - Cruise stage (end of FY)	4.4	4.5	5.1	5.6	6.2	6.3	6.3	6.3	6.4	7.1	8.0	8.0	9.0	9.0	9.0
TRL realized - Lander stage (end of FY)	3.0	5.6	7.0	7.7	7.7	8.0	8.0	8.0	8.0	8.0	8.0	8.0	9.0	9.0	9.0
TRL combined (min)	3.0	4.5	5.1	5.6	6.2	6.3	6.3	6.3	6.4	7.1	8.0	8.0	9.0	9.0	9.0
Ready for launch? (both TRL>=8)	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1
Ready-for-launch & FY>=target?	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1
Cruise Phase Complete?	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1
Labor cost	$12,500,000	$8,463,611	$22,500,000	$27,500,000	$19,341,054	$37,500,000	$30,000,000	$13,385,778	$10,000,000	$3,738,420	$7,500,000	$7,500,000	$2,177,108	$0	$0
Non-labor + overhead cost	$77,235,772	$66,927,223	$85,583,147	$98,050,965	$85,911,742	$115,391,185	$134,644,198	$76,771,557	$77,023,335	$57,476,840	$66,638,572	$64,417,327	$53,585,128	$0	$0
Total annual cost (=spend)	$89,735,772	$75,390,834	$108,083,147	$125,550,965	$105,252,795	$152,891,185	$164,644,198	$90,157,335	$87,023,335	$61,215,260	$74,138,572	$71,917,327	$55,762,237	$0	$0
Expected benefit (ops only)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Cashflow (=benefit - cost)	-$89,735,772	-$75,390,834	-$108,083,147	-$125,550,965	-$105,252,795	-$152,891,185	-$164,644,198	-$90,157,335	-$87,023,335	-$61,215,260	-$74,138,572	-$71,917,327	-$55,762,237	$0	$0
DCF (start-of-year timing)	1.0	1.0	1.0	0.9	0.9	0.9	0.9	0.9	0.9	0.8	0.8	0.8	0.8	0.8	0.8
Present value of cashflow	-$89,735,772	-$73,985,117	-$104,090,144	-$118,658,131	-$97,619,573	-$139,159,073	-$147,062,291	-$79,028,143	-$74,858,695	-$51,676,375	-$61,418,949	-$58,467,903	-$44,488,727	$0	$0
Net present value	-$1,140,248,894														
Launch on Target?	0														

Figure 12: Uncertainty NPV Case

A.3 Flexibility NPV (Mitigated) TRL Progression Mission Lifecycle Model

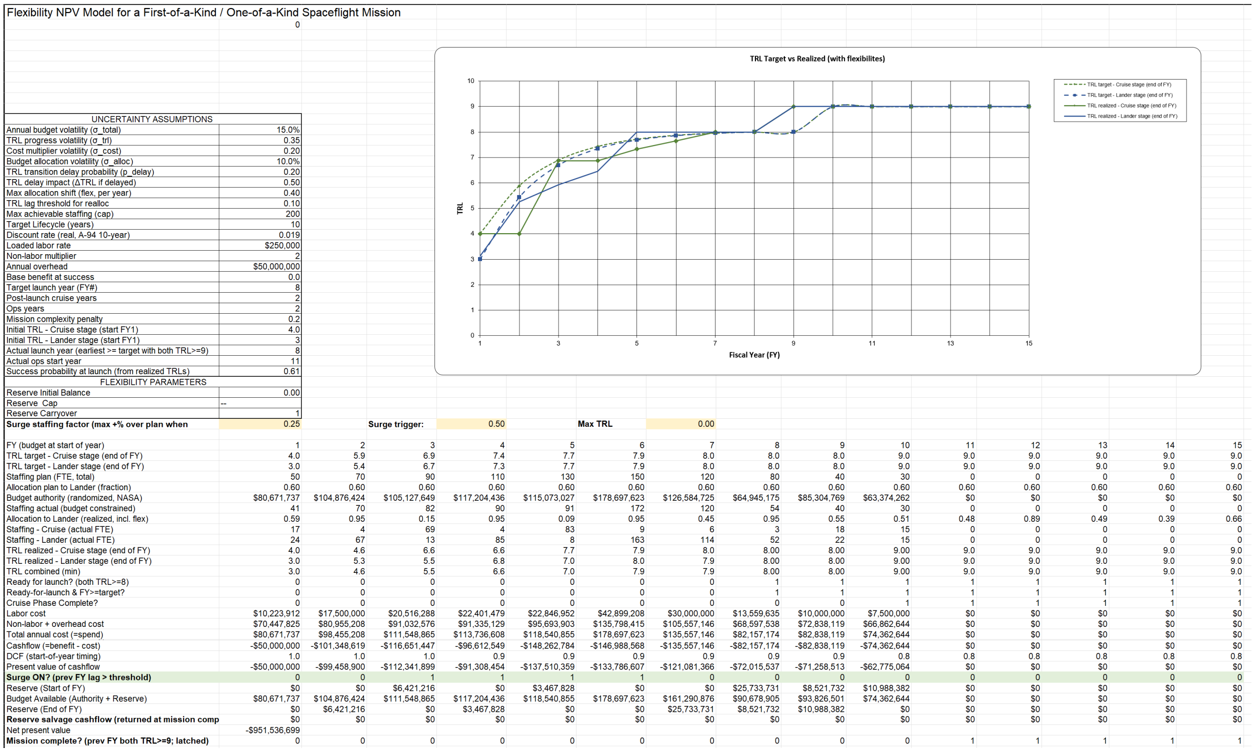

Flexibility NPV Model for a First-of-a-Kind / One-of-a-Kind Spaceflight Mission

0

UNCERTAINTY ASSUMPTIONS	
Annual budget volatility (σ_total)	15.0%
TRL progress volatility (σ_trl)	0.35
Cost multiplier volatility (σ_cost)	0.20
Budget allocation volatility (σ_alloc)	10.0%
TRL transition delay probability (p_delay)	0.20
TRL delay impact (ΔTRL if delayed)	0.50
Max allocation shift (flex, per year)	0.40
TRL lag threshold for realloc	0.10
Max achievable staffing (cap)	200
Target Lifecycle (years)	10
Discount rate (real, A-94 10-year)	0.019
Loaded labor rate	$250,000
Non-labor multiplier	2
Annual overhead	$50,000,000
Base benefit at success	0.0
Target launch year (FY#)	8
Post-launch cruise years	2
Ops years	2
Mission complexity penalty	0.2
Initial TRL - Cruise stage (start FY1)	4.0
Initial TRL - Lander stage (start FY1)	3
Actual launch year (earliest >= target with both TRL>=9)	8
Actual ops start year	11
Success probability at launch (from realized TRLs)	0.61
FLEXIBILITY PARAMETERS	
Reserve Initial Balance	0.00
Reserve Cap	--
Reserve Carryover	1
Surge staffing factor (max +% over plan when	0.25

Surge trigger: 0.50

Max TRL 0.00

FY (budget at start of year)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
TRL target - Cruise stage (end of FY)	4.0	5.9	6.9	7.4	7.7	7.9	8.0	8.0	8.0	9.0	9.0	9.0	9.0	9.0	9.0
TRL target - Lander stage (end of FY)	3.0	5.4	6.7	7.3	7.7	7.9	8.0	8.0	8.0	9.0	9.0	9.0	9.0	9.0	9.0
Staffing plan (FTE, total)	50	70	90	110	130	150	120	80	40	30	0	0	0	0	0
Allocation plan to Lander (fraction)	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60
Budget authority (randomized, NASA)	$80,671,737	$104,876,424	$105,127,649	$117,204,436	$115,073,027	$178,697,623	$126,584,725	$64,945,175	$85,304,769	$63,374,262	$0	$0	$0	$0	$0
Staffing actual (budget constrained)	41	70	82	90	91	172	120	54	40	30	0	0	0	0	0
Allocation to Lander (realized, incl. flex)	0.59	0.95	0.15	0.95	0.09	0.95	0.45	0.95	0.55	0.51	0.48	0.89	0.49	0.39	0.66
Staffing - Cruise (actual FTE)	17	4	69	4	83	9	6	3	18	15	0	0	0	0	0
Staffing - Lander (actual FTE)	24	67	13	85	8	163	114	52	22	15	0	0	0	0	0
TRL realized - Cruise stage (end of FY)	4.0	4.6	6.6	6.6	7.7	7.9	8.0	8.00	8.00	9.00	9.0	9.0	9.0	9.0	9.0
TRL realized - Lander stage (end of FY)	3.0	5.3	5.5	6.8	7.0	8.0	7.9	8.00	8.00	9.00	9.0	9.0	9.0	9.0	9.0
TRL combined (min)	3.0	4.6	5.5	6.6	7.0	7.9	7.9	8.00	8.00	9.00	9.0	9.0	9.0	9.0	9.0
Ready for launch? (both TRL>=8)	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
Ready-for-launch & FY>=target?	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
Cruise Phase Complete?	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1
Labor cost	$10,223,912	$17,500,000	$20,516,288	$22,401,479	$22,846,952	$42,899,208	$30,000,000	$13,559,635	$10,000,000	$7,500,000	$0	$0	$0	$0	$0
Non-labor + overhead cost	$70,447,825	$80,955,208	$91,032,576	$91,335,129	$95,693,903	$135,798,415	$105,557,146	$68,597,538	$72,838,119	$66,862,644	$0	$0	$0	$0	$0
Total annual cost (=spend)	$80,671,737	$98,455,208	$111,548,865	$113,736,608	$118,540,855	$178,697,623	$135,557,146	$82,157,174	$82,838,119	$74,362,644	$0	$0	$0	$0	$0
Cashflow (=benefit - cost)	-$50,000,000	-$101,348,619	-$116,651,447	-$96,612,549	-$148,262,784	-$146,988,568	-$135,557,146	-$82,157,174	-$82,838,119	-$74,362,644	$0	$0	$0	$0	$0
DCF (start-of-year timing)	1.0	1.0	1.0	0.9	0.9	0.9	0.9	0.9	0.9	0.8	0.8	0.8	0.8	0.8	0.8
Present value of cashflow	-$50,000,000	-$99,458,900	-$112,341,899	-$91,308,454	-$137,510,359	-$133,786,607	-$121,081,366	-$72,015,537	-$71,258,513	-$62,775,064	$0	$0	$0	$0	$0
Surge ON? (prev FY lag > threshold)	0	0	1	1	1	1	0	0	0	0	0	0	0	0	0
Reserve (Start of FY)	$0	$0	$6,421,216	$0	$3,467,828	$0	$0	$25,733,731	$8,521,732	$10,988,382	$0	$0	$0	$0	$0
Budget Available (Authority + Reserve)	$80,671,737	$104,876,424	$111,548,865	$117,204,436	$118,540,855	$178,697,623	$161,290,876	$90,678,905	$93,826,501	$74,362,644	$0	$0	$0	$0	$0
Reserve (End of FY)	$0	$6,421,216	$0	$3,467,828	$0	$0	$25,733,731	$8,521,732	$10,988,382	$0	$0	$0	$0	$0	$0
Reserve salvage cashflow (returned at mission comp	$0	$0	$0	$0	$0	$0	$0	$0	$0	$0	$0	$0	$0	$0	$0
Net present value	-$951,536,699														
Mission complete? (prev FY both TRL>=9; latched)	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1

Figure 13: Flexibility NPV Case

Works Cited

References

[1] Richard de Neufville and Stefan Scholtes. *Flexibility in Engineering Design.* MIT Press, Cambridge, MA, 2011.

[2] William M. Kimmel. Technology readiness assessment: Best practices guide. `https://ntrs.nasa.gov/api/citations/20205003605/downloads/%20SP-20205003605%20TRA%20BP%20Guide%20FINAL.pdf`, 2020. NASA Office of the Chief Engineer / NASA Office of the Chief Technologist.

[3] MIT System Design and Management. Garage template 2026 version, 2026. Spreadsheet model template used as the basis for the uncertainty and flexibility analysis.

[4] NASA Jet Propulsion Laboratory. Illustration of spacecraft sections, n.d. Image credit: NASA/JPL.

[5] Office of Management and Budget. Circular a-94: Guidelines and discount rates for benefit-cost analysis of federal programs. `https://www.whitehouse.gov/wp-content/uploads/2025/12/a094.pdf`, 2025. Official OMB circular text.