Reliability and Cost Modeling of Reusable Launch Vehicles

Predicting, Preventing and Mitigating the Cost of Failure

Gonçalo Vera-Cruz

Faculty of Aerospace Engineering

Reliability and Cost Modeling of Reusable Launch Vehicles

Predicting, Preventing and Mitigating the Cost of Failure

by

Gonçalo Vera-Cruz

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday June 28, 2022 at 10:00 AM.

Student number:	5126304	
Project duration:	September 6, 2021 – May 24, 2022	
Thesis committee:	Dr. J. Guo,	TU Delft, Chair
	ir. B. T. C. Zandbergen,	TU Delft, Supervisor
	ir. N. T. Drenthe,	ESA, Supervisor
	ir. R. Noomen ,	TU Delft, External Examiner

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Preface

Almost six years ago, a long night started. It wasn't a cold, lonely night. On the contrary, it was one of those exciting nights, teeming with promise and possibility. It was also by no means a restful night, but it was populated with big dreams, and matching dreamers to achieve them.

It all started with the first dream, which was studying Aerospace Engineering. It was so massive and important at the time, and now seems so far away. Then came making a rocket that could reach the stars, and that was achieved, with a lot of dedication from many other dreamers. The dream changed, and sleep took me to distant lands, where I got to know a different reality and learn from the best. Another dream was lived, and I got to join ESA and work together with people coming from different countries towards the same dream. It now culminates with this present report, where I add my little contribution to the fabric of science, and hope it is useful to others like me in the future.

Like every other big night, there were also nightmares. A pandemic that stopped the dream reel dead in its tracks. And a big dreamer that was left behind and didn't make it to the end of the night.

This night cannot end however without thanking all the people who made it the fantastic journey that it was:

To my indispensable dream guides:

Barry Zandbergen, for his constant encouragement, tireless guidance, and abundant transfer of knowledge; who whipped me into shape from the get-go and managed to make a decent researcher out of me.

And Nigel Drenthe, for his boundless enthusiasm, selfless availability, and stirring creativity; for behaving like a true mentor, and helping me navigate the deep waters I decided to wander into.

And because any good dream necessarily has an eclectic and unique set of characters, I thank the cast that populated my adventures:

The family, friends and colleagues from Lisbon, Delft, Darmstadt and Leiden; the professors, supervisors and mentors; a particular thank you to Michel van Pelt, for humoring my proposal of joining the cost section for my Master's project. The roommates who shared this coming of age journey with me. And all those I had the pleasure to know and dream with.

Finally, a big and sincere thank you to my parents, Paulo and Carla Vera-Cruz, my first teachers, for allowing and encouraging me to chase these silly dreams. To my brother Bernardo, who manages to be an example to me, despite being my youngest. And Babette, for her unconditional support, sacrifice, and love, throughout the best dreams and worst nightmares.

Now that dawn approaches, I feel nothing but grateful that my dreams were a reality, and look forward to keep daydreaming my life away.

Gonçalo Vera-Cruz
Delft, May 2022

Summary

The renaissance of reusable space launch vehicles that is being witnessed in the current space industry demands a new cost estimating approach, which takes into account not only the time dependency in the life-cycle brought upon by reusability, but also the complementary considerations of minimizing cost and reliability.

To achieve this end, a methodology was developed which allows to estimate both cost and reliability, and in a trade-space test various configurations at the development and operating phases, in order to achieve an optimum configuration which minimizes cost, maximizes reliability, and fulfills the design requirements.

The development, manufacturing and operations costs of reusable launch vehicles were obtained using the Theoretical First Unit equivalence method. This includes the particular cost features of reusability, including the recovery hardware, the retrieval of the reusable equipment and its refurbishment. Furthermore, a model for the failure cost was developed, including: Flight replacement, Insurance Penalties, Failure Investigation, Modifications and Downtime.

A reliability model is implemented in parallel to the cost model, which allows to obtain a flight-dependent reliability estimate, taking into account the possible reliability increase methods, such as redundancy, derating, engine-out capability, and extensive testing. This is achieved using publicly available test and operational data and using non-parametric (Kaplan-Meier) as well as parametric (Exponential and Weibull Maximum Likelihood Estimates) techniques in order to obtain life-cycle reliability.

These two models coalesce into a trade-space where it is possible to tune the launcher configuration in order to meet its requirements, while minimizing cost and/or maximizing reliability. At the design level, the features of different configurations and testing programs can be compared, in order to achieve the lowest cost possible while attaining the necessary reliability targets. On the other hand, at the operation level, the reliability figures can be used in conjunction with the cost per flight in order to obtain the expected cost of failure and value throughout the life-cycle of the launch vehicle. This allows the designer to forecast the potential financial losses throughout the operating life of the vehicle, and look to minimize these losses while securing proper funding to cover them, which is especially critical in the first years of operations.

Two use-cases are presented. Firstly, the model is applied to the Falcon 9 vehicle. Verification of the cost figures shows that the error of the methodology is under 30%. A projection of the life-cycle costs is obtained, as well as the expected cost of failure and expected value. Recommendations are given in order to mitigate losses and advise the planning of the life-cycle of the fleet and its reuses. It was found that at an initial stage, where cost reduction from learning and reliability growth haven't yet taken hold, flights have a negative expected value, meaning that the launch provider can expect to lose money with these flights. Insurance options are investigated in order to mitigate these losses.

Secondly, a study is done based on the Ariane 6 publicly available data, in order to study the impact of varying the number of engines in its first stage as well as engine commonality with the upper stage. It was found that for an expendable vehicle such as the Ariane 6, a single-engine configuration provides the most value, as the manufacturing costs of producing several engines per flight does not compensate eventual lower development costs. On the other hand, it was found that for a reusable vehicle, the stricter the reliability requirements, the more a multi-engine configuration is advantageous. This is due to the lower development costs combined with the replacement of the high cost of manufacturing engines with refurbishment costs, with an optimum for a configuration of 4-5 engines.

Contents

List of Figures

List of Tables

List of Acronyms

aces	Advanced Cost Estimating System.
AMSAA	Army Materiel Systems Analysis Activity.
CBS	Cost Breakdown Structure.
CEMs	Cost Estimating Methods.
CERs	Cost Estimating Relationships.
CPI	Consumer Price Index.
CSM	Command and Service Module.
DLR	Deutsches Zentrum für Luft- und Raumfahrt.
DRLS	Down-Range Landing at Sea.
ELV	Expendable Launch Vehicle.
ESA	European Space Agency.
ESOC	European Space Operations Centre.
F9	Falcon 9.
FMECA	Failure mode, effects, and criticality analysis.
FRT	First Stage Reusability Tool.
FTA	Fault-Tree Analysis.
GTO	Geostationary Transfer Orbit.
IAC	In-Air Capture.
ISRO	Indian Space Research Organisation.
KME	Kaplan-Meier Estimator.
LEO	Low-Earth Orbit.
LSE	Least Squares Estimator.
LV	Launch Vehicle.
LVCM	Launch Vehicle Cost Model.
MAP	Maximum-a-posteriori.
MLE	Maximum Likelihood Estimator.
NAFCOM	NASA/Air Force Cost Model.
NASA	National Aeronautics and Space Administration.
OCM	Operations Cost Model.
OECD	Organisation for Economic Co-operation and Development.
PCEC	Project Cost Estimating Capability.
RBD	Reliability Block Diagram.

RFG	Re-Fight Guarantee.
RLV	Reusable Launch Vehicle.
RTLS	Return to Launch Site.
SEER	Systems Evaluation and Estimation of Resources.
SOLSTICE	Small Orbital Launch Systems: a Tentative Initial Cost Estimate.
SPICE	Standard Parametric Information for Cost Engineering.
SRB	Solid Rocket Booster.
SRM	Solid Rocket Motor.
SSCM	Small Satellite Cost Model.
SSME	Space Shuttle Main Engine.
STS	Space Transportation System.
TPS	Thermal Protection System.
US	United States.
VTVL	Vertical Take-off, Vertical Landing.
WBS	Work Breakdown Structure.

Nomenclature

α	Reliability Growth Factor
β	Common-Cause Failure Fraction
χ	propellant reliability factor
η	Derating Factor
K	Uprated design constant
κ	Shape parameter
λ	Scale parameter
ω	Weighting Factor
ϕ	fraction of failures affected by operational derating
σ	Percentage of Work Subcontracted
θ	Generic Parameter
AF	Adjustment Factor
$arg\ max$	Argument of the Maximum
$arg\ min$	Argument of the Minimum
avg	Average
b	learning exponent
C	Cost
c	Specific Cost
C_F	Expected Cost of Failure
C_f	Failure Cost
c_p	Cost Reduction Factor
CF	Catastrophic Failure fraction
CI	Confidence Interval
CpF	Cost per Flight
d	failure
DD	Design and Development
DEV	Development Cost
DM	Development Model
DOC	Direct Operations Cost
DT	Downtime

EM Engineering Model

ENG Engineering

Eng Engine

EO Engine-Out

EUR Euro

EV Expected Value

Exp Exponential

f probability density function

f_4 Learning Factor

f_8 Country Productivity Factor

f_c Assembly and Integration Factor

f_v Vehicle Type Factor

f_{11} Commercialization Factor

f_{dt} Downtime Factor

FM Flight and Mission Operations Cost

$FM1$ Flight Model

FU Flight Unit

FY Fiscal Year

$GLOW$ Gross Lift-Off Weight

GND Ground Operations Cost

h hazard function

HW Hardware

$I\&T$ Integration and Testing

ins insurance rate

IOC Indirect Operations Cost

K Failure rate of first flight

L Likelihood

L_c, L_d Learning Curves

LCC Life-Cycle Cost

LpY Launch Rate

M Mass

M/PA Management and Product Assurance

$M/PA\%$ Management and Product Assurance Percentage

$MAIT$ Manufacturing, Assembly, Integration and Testing

MAN	Manufacturing Cost
MEQ	Mission Equivalence parameter
MOD	Modification Cost
N	Number of Stages
n	number of units
O/F	Mixture Ratio
OPS	Operations Cost
ox	oxidizer
P	Probability
p	learning factor
PFM	Proto-Flight Model
PO	Project Office
PpF	Price per Flight
$pres$	pressurant
$PROP$	Propellant Cost
Q	Unreliability
Q_N	Vehicle Complexity Factor
QM	Qualification Model
R	Reliability
r	Reusable
R_L	Landing Reliability
R'_L	Landing Reliability Extrapolated from Operational History
REC	Recovery Cost
REF	Refurbishment
Rel	Reliability
REP	Replacement Cost
res	Residuals
S	Survivor Function
$S1$	First Stage
$S2$	Second Stage
SS	Sum of Squares
STH	System Test Hardware
t	time
$T1$	Theoretical First Unit

tpm	Tests per month
tpw	Tests per week
TRL	Technology Readiness Level
U	Outcome of Event
USD	United States Dollar
$V2.1$	Vulcain 2.1
w	exponent for design uprating
WB	Weibull
Wh	Work-hour
WYr	Work-Year
$z_{\alpha/2}$	α-th Quantile of Standard Normal Distribution

1 Introduction

Looking at the current landscape in the launch vehicle industry, one big trend can be identified: Reusability. This technology centers around the recovery of part or the totality of the launcher, which is then refurbished and re-launched [192].

Throughout the world there are both private and public initiatives concerned with launcher reusability: In the American scene, the private sector dominates, with Blue Origin's New Shepard becoming the first Vertical Take-off, Vertical Landing (VTVL) launcher to reach space and execute a propulsive landing [2]. After that, SpaceX's Falcon 9 (F9) best exemplifies the potential of reusable launchers, with 151 total launches, 111 landings and 91 re-flown launchers as of the moment of writing [4]. Other SpaceX initiatives in reusability include the Dragon capsule, the Falcon Heavy launcher and the development of the fully reusable launcher Starship [11]. Finally Rocketlab's Electron is a small launcher which has seen its first stage being recovered from the ocean, and there are plans for mid-air retrieval [3]. In Europe, the European Space Agency (ESA) has several initiatives: its uncrewed laboratory Space Rider [10] has its first mission scheduled for 2023, the same year planned for the in-flight demonstration of a prototype reusable rocket first stage called Themis [12]. Russia has a partially reusable launcher in development called Amur [1]. The Indian Space Research Organisation (ISRO) is also aiming at launcher reusability [7]. Finally, China is looking to implement reusability on its new generation of launchers [9].

Reusability of launch vehicles is however not a novel concept, the most emblematic example being the Space Transportation System (STS), commonly referred to as the Space Shuttle, which was in operation from 1981 to 2011. Despite its popularity, the unexpected lower performance and higher costs [183], combined with other factors, lead to its eventual retiring.

In spite of their apparent popularity and ubiquity, is there really a case for preferring a Reusable Launch Vehicle (RLV), compared to an Expendable Launch Vehicle (ELV)? Answering this question requires demonstrating that there is a cost benefit in employing these vehicles. Boiling this further down, the reusable solution only makes sense if the cost decrease achieved by reusing part or the totality of the vehicle is not supplanted by the extra costs incurred from higher development, production and operations costs, and the added effort involved in recovering and refurbishing the stage [129, 154, 165, 184].

The fact is that at the time of writing this document, SpaceX is the *de facto* leader in the launcher market, offering the flight on a new launcher for $67 million in 2022 [66] and on a reused one for around $50 million [31]. This contrasts with the European offer with Ariane 5 at around $175 million [39].

So what characterizes this type of launchers? Three key words are typically used together by these aforementioned companies, in relation to their launchers' development: Reusability, cost, and reliability [11, 49]. Reliability appears here because of the necessity of having the launcher survive for the subsequent flights. This means that RLVs need to typically be as or more reliable than their ELV counterparts [129]. This is difficult, as the inherently higher complexity of the reusable configuration makes it necessarily less reliable, contrary to the popular misconception [122].

Work can be found in literature on the intersection and interactions between cost and reliability in the space propulsion and launcher sectors [84, 125, 129, 132, 132, 149]. However, a few research gaps were identified:

The literature review showed that even though launch vehicle cost estimation is a process that has been realized since the beginning of space exploration, and that there have been efforts in the past to find a link between the cost and reliability of launch vehicles, these have been mainly applied to the realm of expendable launch vehicles, or reusable launch vehicles which don't represent the current predominance of VTVL vehicles over space-planes [129, 132, 149]. Furthermore, little attention seems to be given to the effect of failure on cost [129, 150]. Finally, there seems to be a gap in how decision-making is affected by these considerations, in a value-based process [123, 124, 182, 187].

The current trend noticeable in industry of a move into reusable launch vehicles, which stems from an observed potential benefit at cost level [86], results in a pressing demand in the near future for better and more accurate costing methodologies for these types of vehicles [155, 165]. Weaknesses already identified in the expendable case [99], such as the estimation of operation costs, combined with the new challenges brought about by the new reusable technologies, such as the costing of the launcher's retrieval and refurbishment [86, 155, 165], constitute a gap in the current technical knowledge in the cost domain.

On the other hand, failure is an important factor, as the nominal operation of an RLV fleet becomes dependent on the survival of the reusable hardware and is exposed to the financial losses caused by aging and consequent declining reliability [129, 150]. Furthermore, the importance of considering reliability as a parameter which influences the cost of life-time flights has been shown [132], but needs to be adapted to the modern reusability landscape.

Investigating these gaps and providing a solution will equip designers with better tools to make the right decisions not only in early phases at conceptual and early design level, but also throughout the administration of the launch vehicle fleet.

To that end, the following research question is proposed:

Is it possible to estimate the life-cycle costs and expected value of a reusable launch vehicle, taking into account the expected cost of failure?

To positively answer this question, a figure for the expected cost of failure will have to be obtained. On top of that, incorporating failure into the costing of the launch vehicle will have to result in a cost figure which reflects both the uncertainty and effects of a failure, on the final price tag of the flight. The criterion which informs as to whether the research question has been adequately answered, with a sufficient degree of accuracy, would depend on having data regarding cost, reliability and failure. If one would consider a reusable launch system, operated it for a large number of flights, and registered the development, manufacturing and operating costs, the reliability evolution, and the costs coming from failures, then this could be compared against the model developed, to determine its accuracy.

In order to guide the way to answer this question, six sub-questions are derived, therefore modeling the process to address the main research question:

1. ***How to estimate the cost of reusable launch vehicle hardware?*** - Considering that the cost of ELVs is appropriately addressed by SOLSTICE [99], what needs to be added in order to model the cost of the recovery hardware currently in use for modern RLVs.

2. ***How to estimate the cost of operations of a reusable launch vehicle?*** - This entails identifying the differences between ELV and RLV operations, and addressing how to cost them, including refurbishment and retrieval operations.

3. ***How to model the reliability of a reusable launch vehicle throughout its life-time?*** - The modeling of ELV reliability and reliability growth is addressed by Krevor [132] and Martino [149]. However, in the RLV case, it is necessary to model how the reliability of a vehicle varies with reuse.

4. ***How to implement and estimate the cost of reliability increase strategies*** - This question seeks to determine what reliability increase techniques are applied to launch vehicles, as well as how to determine their cost, particularly when it comes to testing.

5. ***How does failure factor into the life-cycle cost of a reusable launch vehicle?*** - This question is meant to reflect on the intersection between cost and reliability throughout the life-cycle of the vehicle. The main objective is learning the impact of the expected cost of failure on the total life-cycle cost as well as the cost increase due to failure per launch vehicle over its life-cycle.

6. ***How does the expected value of a reusable launch vehicle vary with reuse?*** - This question is concerned with the decision-making part of assessing a solution, and how the viability changes throughout the life-cycle of the vehicle.

The TU Delft research in the field of cost and reusable launch vehicles serves as a foundation for this project. ”SOLSTICE: Small Orbital Launch Systems, a Tentative Initial Cost Estimate” by Nigel Drenthe [99] provides a framework to work in the cost domain, while ”Reusable Rocket Upper Stage: Development of a Multidisciplinary Design Optimization Tool to Determine the Feasibility of Upper Stage Reusability” by Lars Pepermans [154], ”Launch Vehicle First Stage Reusability: a study to compare different recovery options for a reusable launch vehicle”, by Mark Rozemeijer [165], and ”Design and Optimization of a Small Reusable Launch Vehicle Using Vertical Landing Techniques”, by Stephane Contant [86] provide a technical background, especially in the recovery aspect of the reusable launch vehicles. This research project will look into complementing the work being done at TU Delft in the field of launch vehicle modeling, by addressing specific issues related to reusable launch vehicles in the cost domain, and for the first time incorporating reliability, failure and expected value considerations.

This research resulted from a partnership between TU Delft and the Cost Engineering Section at the European Space Agency (TEC-SYC). The aim of this collaboration was to assess the impact of launch vehicle reliability in the cost estimating process for small commercial launchers, initiated in 2016 with SOLSTICE [99]. The purpose of the model developed is to assess the trade-offs between cost and reliability, both in the design phase and in the operational phase, at the early design stage of future expendable and reusable launch vehicles.

The present report is divided into seven chapters, of which this introduction is the first. The second chapter concerns the definition and scope of the problem. Firstly, it looks to contextualize the reader in the current landscape of reusable launch vehicle market and trends Then, it introduces cost engineering, in the broad context of space systems engineering, and as it has been applied to reusable launch vehicles. Finally, it looks at reliability engineering, presenting the current situation in terms of empirical launch success rate, looking at past works that attempted to model reliability of launch vehicles, and the recommendations given for future work applied to reusable launch vehicles. The third chapter continues the area of reliability modeling, by summarizing the main building blocks used and detailing the reliability methodology used in this thesis. The fourth chapter describes the cost methodology used in this thesis, and contains the new cost models developed in this research, including the cost of testing, cost of failure, and expected value. The fifth and sixth chapters consist of two different applications which are used to demonstrate the usefulness of considering a linked model of cost and reliability. The first one is concerned with the premier reference for reusable launch vehicle modeling considered in this research, the Falcon 9 vehicle. In this chapter, the cost and reliability models are applied and its results are verified. Furthermore, the failure and value models provide insight into the effect of failure on the life-cycle cost of the vehicle. On top of that, the expected value model is applied in order to advise an insurance strategy throughout the life-time of the vehicle. The final application, presented in chapter six, consists of a study of the impact of a single-engine first stage configuration versus multi-engine configurations on the cost and reliability of a launch vehicle. The Ariane 6 launcher is used as an example, and on top of this, an assessment on the effect of engine commonality with the upper stage and reusability capabilities are considered. In the seventh chapter, the conclusions coming from this body of work are summarized, while determining if the research questions proposed were answered. Finally, recommendations for further research are given.

2
State of The Art

The purpose of this chapter is to obtain a definition and overview of the landscape of reusable launch vehicles. The main companies and their current and future launchers are identified and described. This provides a context and background information, as well as a general sense of direction in the industry, which is useful to orient this research's modeling objectives. At the same time, the scope of the study is defined.

Furthermore, the cost engineering discipline is here introduced, and its main methods are analyzed, so that the best approach to the problem at hand is determined. At the same time, previous work related to cost estimating of reusable launch vehicles, is referenced, and their strengths and shortcomings are assessed.

Finally, reliability is introduced, giving insight into the main trends in the industry and expectations for the future reusable launch vehicles. Previous works concerned with the reliability modeling of launch vehicles are assessed, as well as their recommendations for further work, in order to clarify what needs to be improved to address the research questions in this project.

2.1. Reusability

This section is concerned with the reusability capability of launch vehicles. It seeks to define reusable launch vehicles in the scope of this study, and gives insight into the future of this technology.

2.1.1. Defining Reusable Launch Vehicle

A Launch Vehicle (LV) is a transportation system designed to launch a payload or crew safely, intact and accurately, from the surface of a celestial body to a target orbit about said body [192]. This study shall focus on the conventionally cylindrical shaped, rocket-propelled LVs with orbital capabilities. This is intended in order to limit the scope of the study, since the majority of the existing cost and optimization tools available are developed to be applied to this configuration.

A Reusable Launch Vehicle (RLV) is defined in the scope of this study as a launch system with the capability of having part or the totality of its composing stages recovered with the purpose of being re-flown. This differs from an Expendable Launch Vehicle (ELV), the most ubiquitous space launch vehicle type to this day, which has its components discarded after use.

To this day, only 4 orbital reusable launch systems have been flown and reused: The Space Transportation System (STS), commonly known as the Space Shuttle; Boeing's X-37 space-plane, and SpaceX's Falcon-9 and Falcon Heavy launchers. The Space Shuttle Program, carried out by the National Aeronautics and Space Administration (NASA) is now retired, being the costliest United States (US) space program ever undertaken [156], and the X-37 is currently being operated by the United States Space Force, and therefore not serving commercial purposes. None of these launchers are fully reusable.

2.1.2. Reusable Launch Vehicles in the present and near-future

Given the potential advantages of launch vehicle reusability, several different private and public initiatives are currently in the development and testing phase. Some of the most relevant ones are presented here.

SpaceX

SpaceX's fleet currently consists of two different launchers with reusability capabilities:

The Falcon 9 [171] is a two-stage launcher with cargo and crew transport capabilities. It has a payload capacity of 22,800 kg to LEO and 8,300 kg to GTO [66]. It was the first commercial orbital class reusable rocket, with a track record of 151 launches, 111 landings and 91 re-flown rockets at the moment of writing [4]. Its first stage counts with nine Merlin engines and aluminum-lithium alloy tanks, using liquid oxygen and kerosene RP-1 as propellant. It employs VTVL techniques, using the Merlin engine restart capability and four carbon-fiber/aluminum honeycomb composite legs to land vertically, on a drone-ship in the ocean, or in a landing zone close to the launch pad. The second stage is powered by a single Merlin Vacuum Engine and is not recovered. The composite interstage employs pneumatic pushers for stage separation, and is equipped with four hypersonic grid fins that allow for moving the center of pressure of the first stage and orient it during reentry. Cold gas thrusters are used to flip the first stage. Finally, the carbon fiber composite fairing and the Dragon capsule are both recoverable and reusable. The payload fairing has been caught 9 times in total. It is equipped with thrusters and a steerable parafoil, which allow for a controlled descent. This is combined with a fleet of two ships with nets to intercept the fairing during its fall. Currently, the intercept procedure has been canceled and fairings are being recovered after soft splashdown in the ocean. The Dragon capsule [27] has been launched 34 times, and has been re-flown 14 times. It is equipped with 16 Draco thrusters for attitude control, and 8 SuperDraco engines in the launch escape system, in crewed flights. It is recovered using two drogue parachutes and four main parachutes.

The Falcon Heavy launcher [4] is composed of 3 Falcon 9 nine-engine cores. It has a capability of 63,800 kg to LEO and 26,700 to GTO [66]. At the moment of writing, it has accomplished 3 total launches, 7 landings and 4 re-flown stages.

SpaceX is currently developing the Starship launcher [68], composed of the homonym Starship upper stage and the Super Heavy first stage. It is intended to be a fully reusable system for orbital insertion, lunar missions and interplanetary transportation, for crew and cargo. It is expected to have a $>$100 ton payload capacity to LEO. It is powered by the full-flow staged combustion rocket engine Raptor, which uses liquid oxygen and methanol as propellant. The spacecraft Starship lands vertically using the gymbaled Raptor engines, and fins for active control.

Blue Origin

Blue Origin's launcher fleet is made up of New Shepard and New Glenn.

New Shepard [51] is a suborbital launch vehicle, the first one to have taken-off and landed vertically. It has flown past the Kármán Line and landed safely back on Earth five times, at the time of writing. The system is fully reusable, made up of a booster and a capsule. The booster uses a top ring to passively displace the center of pressure and thus help control descent, in combination with four wedge-shaped fins. Eight drag brakes are used to decelerate the vehicle, and the engine BE-3PM is restarted to achieve propulsive landing. The hydraulic actuated aft fins are employed during ascent and descent, and finally the landing gear consists of retractable landing legs. The capsule is recovered with parachutes.

New Glenn [50] is an orbital class launcher in the development phase. It is designed for a performance of 13t to Geostationary Transfer Orbit (GTO) and 45t to Low-Earth Orbit (LEO). It is composed of two stages, where only the first one is reusable. The first stage is powered by seven BE-4 engines, and the upper stage by two BE-3U engines. It uses four forward fins for active attitude control on descent, two wing-like strakes for aerodynamic performance, and the landing gear is composed of six hydraulically actuated legs.

Rocketlab

Rocketlab's only flying launcher is the Electron [3]. It is a smaller launcher compared to the ones already mentioned in this section, being specifically geared towards the small satellite market. The first stage is powered by nine sea-level Rutherford engines and the second stage by a single vacuum Rutherford engine. It has a 300kg to LEO capability. Reusability of the first stage is being planned, with a first stage having been recovered after splashdown in the ocean (descent was controlled with an

RCS system, and decelerated with a drogue and parachute system). Aerial capture with a helicopter is currently in the works, with an incomplete attempt being achieved in May of 2022 [57]. In 2021, a successful recovery after splashdown saw an updated thermal protection system tested, behaving successfully, with the team stating that some components were already suitable for reuse, and the engines remaining in good condition and being prepared for hot fire testing [59].

The company is also developing the Neutron launcher [48], which features a reusable first stage designed to land on an ocean platform and a payload capability of 8 tons to LEO.

Others

The United Launch Alliance presented a concept to recover and reuse the booster engines of the Vulcan Centaur launcher. According to this plan, the detachable module would descend under an inflatable heat shield and be captured by helicopter after parachute deployment.

Virgin Galactic is currently testing their SpaceShipTwo prototype [63], which is a fully reusable, suborbital space-plane intended for space tourism.

India has made attempts at reusability, better illustrated with the flight of the RLV-TD in 2016 [58], a reusable two-stage to orbit vehicle, which takes-off vertically. The upper stage is a space-plane that lands horizontally. In 2021, the purpose of pursuing reusability was again emphasized in the new year's message of the secretary and chairman of ISRO [7].

In October 2020, Roscosmos announced the development of a new launch system, Amur [1]. This is expected to be a methane powered reusable launcher with a launch cost of $22M, wet mass of 360t, LEO capacity of 12.5t and reliability of up to 0.99. It is designed with grid fins to brake in the atmosphere, and landing legs for vertical propulsive landing. Its development is expected to cost 70 billion rubles ($925.75M) and the first launch is planed for 2026.

Airbus has a concept for first stage reusability, called Adeline [14]. This would see the engine module being jettisoned on reentry, and land horizontally on a runway using deployable winglets and propellers.

ESA has had several scrapped attempts at reusability. The most recent venture was announced in 2020: The Themis programme [12] is aiming at flying the prototype for a reusable rocket first stage, powered by the Prometheus engine. This is planned to be a VTVL vehicle, using propulsive landing.

Finally, the Chinese company i-Space is developing its launcher Hyperbola-2 [43], a 106t vehicle with a reusable first stage which lands vertically. The launcher expected performance is 1.9t to LEO.

2.1.3. Summary

Given the different RLV solutions being successfully operated or currently in development, a great tendency towards VTVL technologies is noticed. Both the Falcon family of launch vehicles and New Shepard rely on it, being the only vehicles currently flying and being recovered and re-flown. Furthermore, future initiatives such as Themis, Amur, Neutron, also employ VTVL technology. Alternative options, such as parachute recovery or horizontal landing, seem to be less privileged. For this reason, a special attention shall be given in this study to the vertical take-off and vertical landing solution. The model developed will be applicable to different solutions, but priority will be given to the VTVL solution that has been successfully deployed in the market by SpaceX.

2.2. Cost

With the design-to-cost philosophy being employed more ubiquitously in the industry today, cost engineering becomes an important discipline throughout all design phases. Emphasis can be placed on the accuracy and uncertainty of the estimates, as an underestimation can lead to more financial and schedule costs down the line, and overestimation can lead to a loss of contract.

In this section, an overview of the industry's foremost Cost Estimating Methods (CEMs) and cost estimation tools is given, as well as previous works on the field of cost estimation of expendable and reusable launch vehicles.

2.2.1. Cost Estimation Methods

There are three classic cost estimation methodologies widely recognized in industry: Parametric, build-up and analogy [140]. To this, other complementary methods such as expert judgment and rough order of magnitude are added, as well as ”advanced methods”.

Parametric Cost Estimation

Parametric cost estimation relies on Cost Estimating Relationships (CERs), a series of mathematical relationships based on historical data. It is indicated for use in early program phases in a top-down approach, as only base requirements are defined, and the more elaborated and detailed criteria are not yet available [136].

The goal of the CERs is to correlate physical, technical and performance characteristics to program cost. Cost is therefore seen as a function of these characteristics [140]. Deviations from the underlying parameters are allowed for through user defined inputs [179].

The main challenge related to parametric cost estimation is the fact that it is only as good as the underlying data that serves as a base for the CERs [179]. The quality and quantity of data is directly related to the robustness and credibility of the models.

Engineering Build-up Estimation

The engineering build-up approach is a bottom up approach that consists of summing the costs at the lowest level of detail possible, called a Work Breakdown Structure, often coming directly from the experts or engineers that work on the systems [142, 179]. The tangible costs such as material and labor are estimated and an overhead is applied [88, 140, 142, 179].

For this method, a high level of detail needs to be present and all the parameters for systems and sub-systems are known and clearly defined, therefore being unsuitable for projects in an early-phase [88, 179].

The main challenges with the build-up approach is the associated financial and schedule cost, as it is a very resource-intensive process. Furthermore, these characteristics make it so that the estimate has a high inertia and difficulty to adapt to changes [179]. Furthermore, small errors accumulate and are propagated to the final estimate, which can be detrimental, when combined with the high resource demand [140]. It has an inherent high credibility, since cost estimates are highly justifiable from the high level of detail.

Analogy Estimation

Cost estimating through analogy is a process that consists of extrapolating a comparison between the item being costed and another similar item. It is therefore based on a single data point [140]. Due to it being heavily dependent on the analyst's judgment, it is considered to be a subjective method by some authors [142, 179], although its scope of application can range into more objective and bottom-up variations [88].

It is a method with an inherently quick and effective development, as long as there is a suitable element for the comparison. Beyond the similar characteristics, it is necessary to evaluate the similarity in terms of cost drivers, design and production implications, and from that extrapolate cost considerations [140].

The analogy can be further classified as a loose analogy or close analogy, depending on how closely related the comparison elements are.

The main limitation of the analogy method is related to the inherent limitations implied in a comparison, as the historical data point can have a varying degree of applicability to the current situation [88, 140].

Expert Judgment

The classification of Expert Judgment as a cost estimation method is divisive [179], in the sense that its subjective nature results in some criticism and skepticism, however it is widely used in industry, and as support to the other methodologies. It can be employed throughout all phases of project development, and can be useful when data is scarce. It is furthermore not demanding in terms of financial and schedule cost.

Rough Order of Magnitude

Rough Order of Magnitude estimation is a useful early phase method based on knowledge from past experience and readily available industry data [179].

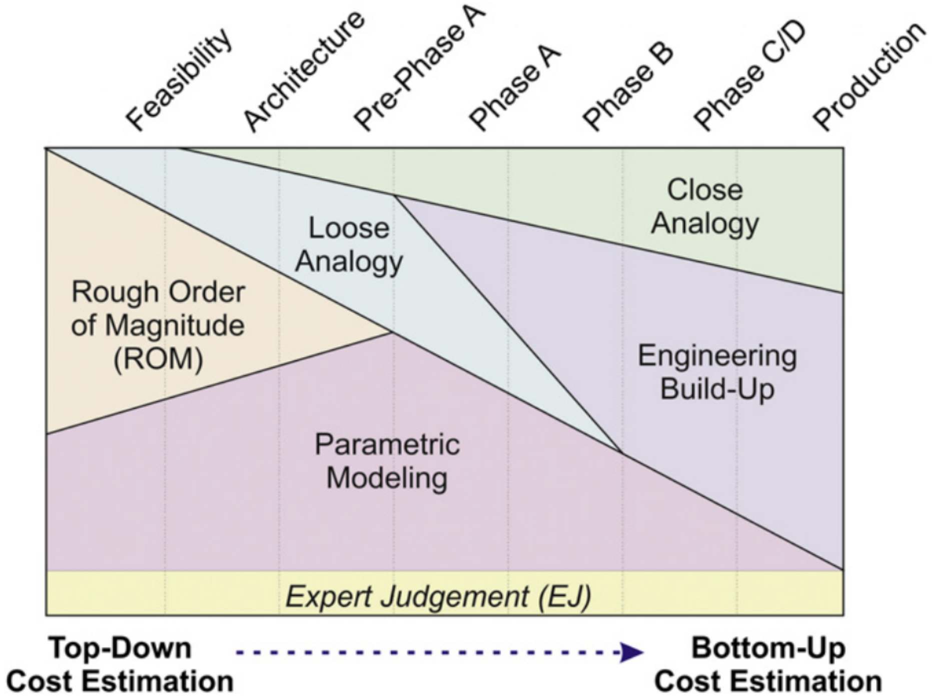

Figure 2.1: Qualitative application of CEMs according to project phase [136].

Advanced Cost Estimating Methods

Despite not being as ubiquitous in literature as the three classic methods and accompanying complements, these methods can provide a new perspective, as well as refine the results obtained.

Feature-based modeling links cost to design features at a system or subsystem level. This logic was brought about with the development of CAD/CAM. As it is a recent technique, inconsistency across the industry comes as a limitation to this method [167].

Fuzzy-logic comes as an answer to the deterministic nature of the classical methods. It looks to include the imprecise and unpredictable nature of the industry in the cost estimates [88].

Neural-networks methods look to train a computer to learn the effects of design features on cost, with historical cases. It has been shown that it can outperform CERs in some circumstances [167]. The main issues with this method are its dependency on a large database, difficulty in adapting to changes in the industry and finally the fact that the underlying relations are not explicit, and difficult to justify, like in the case of CERs.

Summary and Applicability of the Cost Methods

The correct methodology or combination of methods to be applied depends on the phase of the project and the available information, both related to the item to be costed, as well as historical data and similar projects. Furthermore, the knowledge of the experts pertaining to the systems as well as the costing engineers' has influence in the method selection. Figure 2.1 summarizes the applicability of the different cost methods according to project phase. As can be seen from the figure, more than one method is applicable for a certain phase of the project, and so a combination of methods is employed, with expert judgment at the base.

The present research is concerned with early phase projects, when the specific characteristics of the mission and systems aren't yet fully defined. Therefore, top-down approaches to cost modeling are the most adequate ones. Parametric modeling will be the privileged technique, allowing us to leverage the knowledge we have of the system and mission requirements. Analogy methods will also be considered in the cases a parametric approach is not possible. Rough order of magnitude can be employed so as to keep the results and assumptions in check. Finally, expert judgment will be used throughout the project to support the methodology adopted. Throughout this research, while embedded in ESA's Cost Engineering section (TEC-SYC), I had the support of the experts ir. Nigel Drenthe and ir. Michel van

Pelt. Nigel Drenthe was the external supervisor for this thesis project. He is a cost engineer at ESA, having been the developer of the SOLSTICE tool [99], and the author of the article ”Cost estimating of commercial smallsat launch vehicles” [96]. Michel Van Pelt is the head of the Cost Engineering Section at ESA. He has authored multiple publications on topics ranging from cost estimation, systems engineering, space architecture and space exploration [70]. He contributed to this project as an advisor with his expert judgment.

2.2.2. Cost Estimating Tools and Models

Tools and models are used to implement the methodology or methodologies deemed appropriate for the project at hand. The choice of the right tool is dictated by the project's purpose, phase and level of detail available [142].

A number of commercial and government off-the-shelf products exist in the market. There are two main obstacles related to the availability of these tools. Firstly, commercial tools widely used in the aeronautic and space industries, such as Small Satellite Cost Model (SSCM) [137], TruePlanner [54], Advanced Cost Estimating System (aces) [13] and Systems Evaluation and Estimation of Resources (SEER) [60], among others, are not free to use, and require a personal or company/agency-wide license. On the other hand, several government tools, such as NASA/Air Force Cost Model (NAFCOM) [73], Project Cost Estimating Capability (PCEC) [53], Launch Vehicle Cost Model (LVCM) [99], Operations Cost Model (OCM) [103] and Standard Parametric Information for Cost Engineering (SPICE), are developed and owned by government agencies, such as NASA, DoD, or ESA, and are therefore subjected to hard or soft export laws, or their underlying CERs and data are unpublishable, due to confidentiality of their subcontractors.

Two tools stand out, due to their open-source nature and CER/data transparency:

TRANSCOST

The TRANSCOST Space Transportation Cost Model [126, 129] is a launch vehicle costing tool which includes the development, operations and manufacture stages of both expendable and reusable rocket launchers. It is designed for initial conceptual missions, and uses mostly parametric methods based on rudimentary CERs derived from a vehicle and engine database. It is found ubiquitously in research, due to its open and free access.

SOLSTICE

Small Orbital Launch Systems: a Tentative Initial Cost Estimate (SOLSTICE) is the research project of TU Delft alumnus Nigel Drenthe [99]. It estimates the cost of development and manufacturing based on a theoretical first unit equivalent estimate. The first unit cost is obtained parametrically from the mass of the relevant subsystem. The operations cost is estimated with TRANSCOST.

Other models

A number of other models have been developed in order to address the problem of modeling the cost of launch vehicles, usually in order to be integrated into an optimization tool. Several of these make use of publicly available tools, such as TRANSCOST [129], together with select CERs pertaining to specific subsystems or equipment, and bottom-up estimates.

At TU Delft, Contant [86] used SOLSTICE, together with approximations for recovery and refurbishment costs, in order to estimate the cost of a reusable launch vehicle. Pepermans [155] and Rozemeijer [165] used TRANSCOST, as well as a number of CERs available in literature, together with an engineering build-up of certain subsystems, in order to estimate the cost of a reusable upper stage and first stage, respectively.

Elsewhere in literature, some works can be distinguished as being relevant to the present body of work. The most relevant are those by Martino [149], Wertz [149] and Sippel [172].

2.2.3. Cost of Reusable Launch Vehicles

Previous work has shown that depending on the launch rate per annum and the number of reuses, it is possible to achieve a cost reduction by using reusable launch vehicles. Koelle [129] predicts a cost reduction between 35% and 50% using a reusable first stage, or 56% and 70% if also using a reusable upper stage, of a vehicle which is reused 25 times, at a launch rate per year of 15 flights.

Snijders [169] predicts a 50% launch cost reduction by reusing ten times the first stage of a vehicle with a yearly launch rate of 10 flights. On the other hand, Contant [86] finds this model optimistic, finding an optimized configuration which yields a cost reduction of 39.2%. Furthermore, Rozemeijer [165] used Falcon 9 as a case-study, which allowed him to demonstrate that first stage reuse brought a cost decrease close to 30%. Pepermans [155] showed however that reuse of the upper stage resulted in a cost decrease of only 6%.

As noted by Contant and Pepermans, these studies have not taken into account reliability. Increased reliability leads to an increase in cost [129, 132], however, it minimizes the expected costs of failure. Parkinson [150] points out in his article ”The Hidden Costs of Reliability and Failure in Launch Systems” that failure results in added costs beyond the simple re-flight of the satellite. To that end, he also mentions how the cost of failure differs for both expendable and reusable launch vehicles, starting from their potential lower reliability, to the different impact that failure has on their respective programs.

A complete assessment of the potential cost benefits brought about by reusability will necessarily need to consider the expected cost of failure and therefore the vehicle’s reliability. Koelle [129] predicts that the cost of a launch failure could total 2 to 3 times the cost of the flight, whereas Parkinson [150] predicts that this cost can be 3 to 5 times that.

2.3. Reliability

Between 1957 and 2010, counting 4038 launches and 366 catastrophic failures, the overall success rate of launch vehicles was 91.7% [129]. Despite the increasing launch rate per year, the yearly success rate has not increased, with 2020 having the worst success rate in the last 15 years [80]. This is attributed to the comparatively high number of new launch systems in development and flying for the first time, compared to the number of more established and mature ones.

In 2021, the expendable Ariane 5 launcher has flown 112 times, with only 2 catastrophic failures, which translates into a success rate of 98.2% [20]. On the other hand, the Block 5 version of the reusable Falcon 9 vehicle has flown 78 times, including flights on reused boosters, with no failures, totaling a success rate to date of 100% [45]. As for recovery, considering the same version, out of 75 attempts, 4 failed, which results in a success rate of 94.7%. An added particularity is that the Falcon 9 flew 31 times in 2021, a much higher launch rate than the ones achieved in the past.

2.3.1. Reliability Modeling

Previous work in the reliability modeling of expendable launch vehicles are extensive. Zwack [196] uses reliability growth methods with fault tree analysis. On the other hand, Krevor [132] and Martino [149] develop cost and reliability models as part of larger optimization tools, using historic data of failure rates at subsystem and equipment level, coupled with reliability analysis techniques like reliability block diagrams and fault-tree analysis, in order to achieve a reliability estimate at of a vehicle or mission.

Krevor advises further work on the field of the cost of failure, as the added losses brought about by mission failure were not considered, and also recommends that the model be extended to reusable launch vehicles, considering the increasing failure rate that results from reusing a part or the totality of the vehicle. Martino, on the other hand, not only suggests the extending of the tool to the reusable cases, but also recommends considering improving the reliability model by taking into account the qualification process of the engine.

2.3.2. Reliability of Reusable Launch Vehicles

According to Koelle [126], it is expected that the catastrophic failure rate of RLVs is significantly lower than that of ELVs. The reasons for this are increased redundancy, higher safety margins, multi-engine failure capability, possible use of integrated health control systems, the capability of landing in case of emergency and the possibility of testing before operational use.

Parkinson [150] on the other hand postulates that since the reliability of components of reusable launch vehicles compared to the ones used on expendable ones is not expected to increase, given the added complexity of reusable systems, the probability of failure might actually increase, discounting the possibility of having abort modes.

In an article [178], Thiokol Propulsion test the assertion that increased complexity has a direct correlation to reduced reliability, using the FMEA for the STS SSME, coming to the conclusion that increased complexity results in reduced reliability.

The conclusion that can be drawn from this is that given the inherent mission and purpose of RLVs is to be cost effective after a number of launches, then the reliability requirement will be necessarily more stringent, with the final reliability of RLVs outperforming that of ELVs and striving to be comparable to the figures seen in the aviation industry. However, the added complexity that these systems display will have as a consequence a penalty on the reliability, and so designers have to come up with strategies to increase the reliability to the necessary standard, while balancing the cost vs performance variables that tie into these strategies and their result in the reliability.

2.4. Conclusion

Across the public and private sector, it is visible that the trend is shifting towards launch vehicle reusability, with SpaceX leading the sector and acting as an example for a wide range of companies, such as Rocketlab, and governmental agencies around the world investing in this technology, from NASA and ESA to ISRO and Roscosmos. The low price boasted by the Falcon 9 launcher and its reported costs represent a clear benefit, compared to the expendable solutions offered by the competition.

On the cost side, the main engineering cost estimating methods were surveyed, and it was concluded that at the early stage phases of development, in which this research is interested, parametric modeling is most appropriate method. This is supported by analogy and expert judgment. In terms of previous work in estimating launch vehicle costs, multiple attempts have made, but it was found that some parts were lacking when it comes to reusable launch vehicles, specifically the cost of the recovery hardware, its refurbishment, and recovery costs.

On the other hand, it was found that there are established works which outline a methodology for reliability estimating, but in the field of reusable launch vehicles, and how aging affects the launcher, there are some knowledge gaps that can be addressed.

Finally, past works have linked cost to reliability, especially when it comes to the cost of added redundancy, however, measuring the cost of failure and the relation between increased reliability in testing and its cost are lacking.

3

Reliability Model

Reliability is defined as the probability that an item will perform its intended function for a specified interval under specified conditions [112].

It is documented how the effort spent during conceptual and early design stages of a system development cycle can affect up to 80% of the total system costs [168]. Following from this, late application of Reliability, Maintainability, and Availability (RMA) is virtually inconsequential, and so for cost effectiveness purposes, a design-to-reliability philosophy must be applied from the earlier design stages, in order to achieve the desired level of safety and affordability.

For this reason, reliability is a focal aspect of this research. The goal is taking reliability into account in the conceptual design phase and how it relates to cost in the development and operating stages. This will provide a complete notion of the reliability costs throughout the complete life-cycle of the launcher as a product, and allow to make the best decisions in order to improve reliability and decrease cost.

In this methodology, by modeling the reliability of the launcher system, how it evolves throughout the life time of the vehicle, and including it in the cost model, it will be possible to estimate the expected cost of failure. This chapter is concerned with the development of the reliability estimation model, which will be later included in the cost model. On top of that, strategies which lead to reliability increase, and how they are modeled are also investigated at the end of the chapter.

The applications of the methodology described in this chapter will be explored in chapters 5 and 6.

3.1. Analysis Objective, Method and Scope

This section looks to specify the objective of reliability modeling in a greater level of detail. On top of that, the scope of the study is defined, and the methodology is outlined.

3.1.1. Objective

The purpose of modeling launcher reliability is two-fold:

Firstly, the main objective of this exercise is to obtain a figure (or figures) for the probability of failure. With it, it is possible to at any point calculate the expected cost of failure, by combining it with the failure cost model, which shall be introduced in section 4.7.

Secondly, reliability modeling allows to evaluate the impact of reliability increase techniques in the probability of failure. Therefore, it is possible to directly assess its influence on the expected cost of failure. Furthermore, it makes it possible to economically gauge whether the investment into increasing reliability is financially worth it.

3.1.2. Methodology

The method proposed in this methodology for estimating reliability at component and subsystem level is using a lifetime distribution. This method has been used in the past, to estimate the reliability of expendable launch vehicles [132, 149]. The methodology described will be extended to modeling the effects of aging in the reusable launch vehicle case, according to recommendations found in literature [132]. Furthermore, the reliability of events particular to reusability, such as the recovery of the reusable material, are also modeled. Finally, what distinguishes this research from the previous ones that have modeled reliability and cost, is that instead of optimization, the main aim of this project is to measure the impact of failure on the cost of the launch vehicle.

A limitation of this methodology is its high dependency on both available data and expert judgment. The estimate is sensitive to variability of the parameters of the ruling distribution. In this case, the scarcity of data related to launch vehicle reliability, and particularly of reusable VTVL systems, makes it difficult to apply parameter estimating methods, such as least squares regression method, maximum likelihood estimator, method of moments, Bayesian inference, among others.

The onus of selecting an adequate parametrization for the selected lifetime distribution is therefore put on the side of the user. In the best case scenario, it is assumed that the estimator has a complete knowledge of the system and experimental test data. A FMECA analysis paired with expert judgment makes it possible to obtain adequate parameters for lifetime distributions. On the other hand, curve fitting techniques can be applied to experimental test data in order to find parameters for the lifetime distribution. Otherwise, in the worst case scenario, reliability requirements and performance targets can be used with expert judgment to infer lifetime distribution parameters, and effectively enable the reliability analysis and the life-cycle economic analysis.

In the future, with access to data, it is possible to run a better non-parametric estimation and therefore select a better parametric model similar to the work done in satellite reliability by Bouwmeester at TU Delft [82].

Alternative methodologies to model the reliability degradation effect would include stochastic processes [115], such as Markov chains, Stochastic Petri Nets, Gamma Processes, and Brownian motion with drift. These have extensive application in civil engineering to simulate the deterioration of structures [181]. Stochastic processes have been used to model aircraft component degradation in a recent TU Delft study [134]. These methods are more complex and computationally intensive when compared to the more simple lifetime distribution model [132]. For that reason, the latter method is employed in this study.

The reliability calculation at system level is achieved with Fault-Tree Analysis (FTA). This method has heritage in previous launch vehicle reliability studies [132, 149, 191, 196].

The methodology used in this section is illustrated by figure 3.1. In this chapter the reliability estimating methods, reliability analysis techniques and reliability increase strategies are presented, discussed and selected. In chapter 5, the model is applied to the Falcon 9 vehicle, in order to validate it and demonstrate a use-case in which the results are combined with the cost estimate in order to obtain the expected cost of failure.

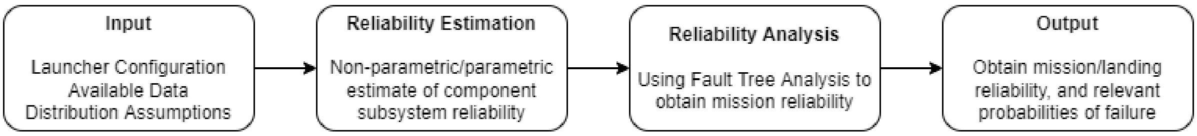

Figure 3.1: Reliability Modeling Methodology.

3.1.3. Scope

This study is concerned with modeling the occurrence of catastrophic failures. These are defined as the failures which result in a loss of the launch vehicle [132]. To that end, reliability is estimated at subsystem level. System and launcher level estimates are obtained with reliability analysis techniques.

Reliability increase techniques are also the object of this study. At subsystem level, redundancy might mitigate catastrophic failure. The propulsion system in specific should allow the conversion of possible catastrophic failures into survivable and benign failures [121, 132]. This is achieved by using a health monitoring system to enable engine-out capability. The possibility of increased and improved testing is also studied.

3.2. Reliability Breakdown Structure

The reliability study begins with breaking the problem into smaller levels of detail, similarly to a Work Breakdown Structure (WBS) in cost estimation [102]. Considering a launch vehicle, reliability can be calculated on multiple indenture levels. At the same level, all elements are of the same complexity.

A typical blueprint for a reliability breakdown structure would be as follows:

- Architecture;
- Element/System;
- Subsystem;
- Component;
- Assembly;
- Part.

Reliability at the level of assembly or parts is usually done in the detailed design phase, when more detailed information of the design is known. For conceptual design, the component level reliability estimate should suffice [132].

Adapted from Krevor [132] the following breakdown is achieved to address the problem in this research:

- Vehicle;
  - Stage;
    - Subsystem;
      - Component;
  - Events;

The vehicle level estimate results from the physical systems estimate (the reliability of the stages), as well as the reliability of the events modeled in the mission, such as firings, separation and landing. The stages are further broken down into subsystems, such as power, avionics and propulsion. It is at this level that most of the reliability estimates are made in this research.

The subsystems can be further broken down into components or equipment. Since the propulsion system is historically the cause for the most failures in launch vehicles [106, 113, 117, 121, 129, 132], it will be in this research modeled at component level, paying particular attention to the engine.

This breakdown will be the input necessary to build reliability block diagrams and fault tree analysis, as will be explored in section 3.4 and applied in chapter 5, section 5.4.5.

3.3. Reliability Estimation Model

Ideally, the reliability estimation at component level is derived from failure data. This mostly pertains to if the outcome of a test was a failure or not, and the time (or cycle, flight, landing, etc.) at which the failure occurs or when the test was stopped. On top of this the failure mechanism is also interesting, for instance for a FMECA analysis. The more extensive and varied the data is, the better the modeling of the system's survival behavior. The estimation can be done in a non-parametric way, using for instance a Kaplan-Meier estimator, or parametric, assuming the component's survival follows a certain distribution, and obtaining the parameters that rule it, for instance using Maximum Likelihood Estimation and/or expert judgment. This type of analysis can be applied at equipment level, all the way to mission level. In the past, this methodology has been applied to a population of different expendable and reusable orbital launch systems, over their total launch history [81]. This methodology can be applied to the operational history of reusable launch vehicle systems, in order to obtain an estimate of its reliability and how it evolves with aging.

3.3.1. Data

Testing or operational data can be classified into two categories: failure data and censored data [135]. A complete data set is made up of only failure data, where the exact time of failure is recorded, as exemplified in figure 3.2a. However, often a test or sample includes survivors. These constitute censored data, as the time of failure is not known. These are called right censored data points, and are represented in figure 3.2b. It is also possible to have left censored data, which is used for items that have failed before the test started, or interval censored, for when there is uncertainty in the time of failure, but the lower and upper bounds are known. These two are less relevant for the problem at

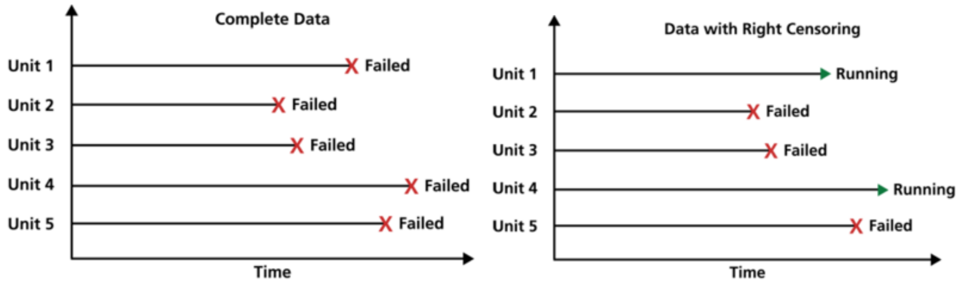

(a) Complete data set. (b) Right Censored data.

Figure 3.2: Complete and right censored data set.s Source: [55]

hand. Furthermore, the observation or testing can be stopped after a certain time has passed, creating Type I censoring, or after a certain number of failures has been achieved, which constitutes Type II censoring. When observing for instance the operation life of a launcher, Type I censoring is used.

It is not possible to perform analysis with no failure data or only right-censored data, so it is extremely important to have a rich and diversified data set. This lead to difficulties when looking at extremely clean track records of very reliable launch vehicles such as the Falcon 9. This is compounded even more due to SpaceX being a private company, resulting in there not being a lot of publicly information available about both their testing programs, and the failures that might have happened in flight but were inconsequential for achieving their mission or landing the first stage.

3.3.2. Non-parametric Estimation

The Kaplan-Meier Estimator (KME) provides an expectation for the survival function over time [101]:

$$\hat{S}(t) = \prod_{i:t_i<t} (1 - \frac{d_i}{n_i}) \tag{3.1}$$

In this method, the time t_i is updated whenever a number of failures d_i occur, having n_i operational units at risk (that is, the total number of units minus the failed and censored ones). The variance and confidence intervals are obtained with the Greenwood method with normal approximation [118]:

$$var(\hat{S}(t)) = \hat{S}(t)^2 \sum_{i:t_i<t} (\frac{d_i}{n_i(n_i - d_i)}) \tag{3.2}$$

$$CI = \hat{S}(t) \pm z_{\alpha/2}\sqrt{var(\hat{S}(t))} \tag{3.3}$$

In this equation, $z_{\alpha/2}$ consists of the α-quantile of the of the normal distribution. For a 95% confidence interval, $\alpha = 0.05$ and $z_{0.025} = 1.96$. An example of the application of a Kaplan-Meier Estimator to a satellite failure database is show in figure 3.3.

The Kaplan-Meier estimator is an exact representation of the failure data, being therefore unbiased and free from assumptions [82]. It is a useful first step in a reliability estimate, as it allows to derive inputs for guesses in parametric estimates, and as a reference to assess goodness-of-fit of a parametric model. This is exemplified in chapter 5, where in section 5.4, non-parametric estimates of the Falcon 9 engine and landing reliabilities are obtained via application of the Kaplan-Meier Estimator, before starting a parametric estimate.

3.3.3. Parametric Estimation

The reliability $R(t)$ at a component or subsystem level expresses the probability that an item with lifetime T is functioning at one particular time t. It can be obtained from the survivor function $S(t)$ of a given failure distribution, which is the probability that an item is functioning at any given time t [135].

$$R(t) = S(t) \tag{3.4}$$

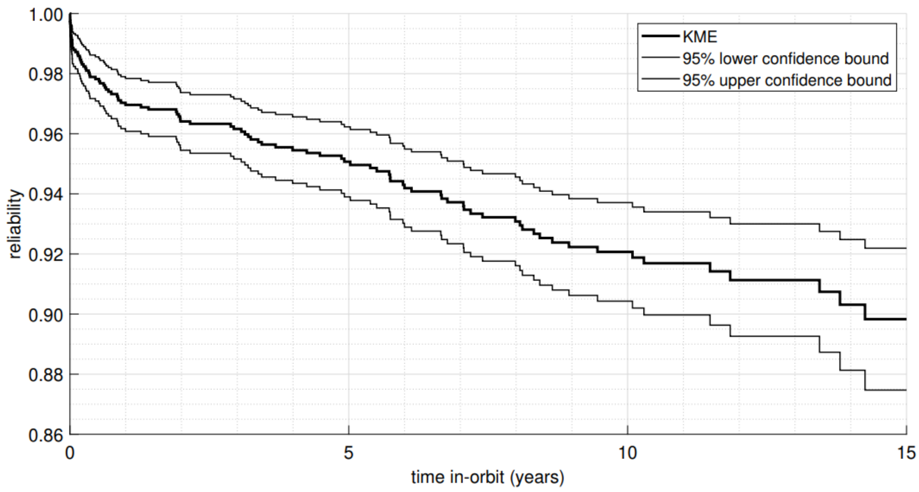

Figure 3.3: Application of a Kaplan-Meier Estimator to a satellite failure database [82].

$$S(t) = P[T \geq t] \quad t \geq 0 \tag{3.5}$$

The survivor function is obtained from the cumulative distribution function of a given distribution as follows:

$$S(t) = 1 - F(t) \tag{3.6}$$

Distributions that are typically used to model the lifetime of a system are called lifetime distributions.

The failure rate of an item is modeled with the hazard function of a given lifetime distribution $h(t)$. The hazard function can be obtained from the distribution's probability density function $f(t)$ (pdf) and survival function [104]:

$$h(t) = \frac{f(t)}{S(t)} \tag{3.7}$$

The two most popular lifetime distributions are the exponential distribution and the Weibull distribution. The exponential distribution is characterized by a having a constant failure rate, whereas the Weibull distribution allows to model a decreasing or increasing failure rate with time.

Other popular lifetime distributions, such as the Gamma, Gompertz, log-logistic and lognormal distributions are discarded from this study, as the Weibull distribution constitutes a flexible distribution that allows to model different scenarios, from immaturity to wear-out [82], and has previously been used in the modeling of engine components [106].

Exponential Distribution

The exponential distribution plays a central role in reliability engineering, as it is the only continuous distribution with a constant hazard function [82, 135]. When considering reliability studies, the hazard function is used to represent the failure rate of a given system. The exponential distribution models therefore a system whose failure rate does not vary with time.

The survivor and hazard functions of the exponential distribution are given by the following equations, and are represented in figure 3.4:

$$S(t) = e^{-\lambda t} \tag{3.8}$$

$$h(t) = \lambda \tag{3.9}$$

Where λ is a positive scale parameter.

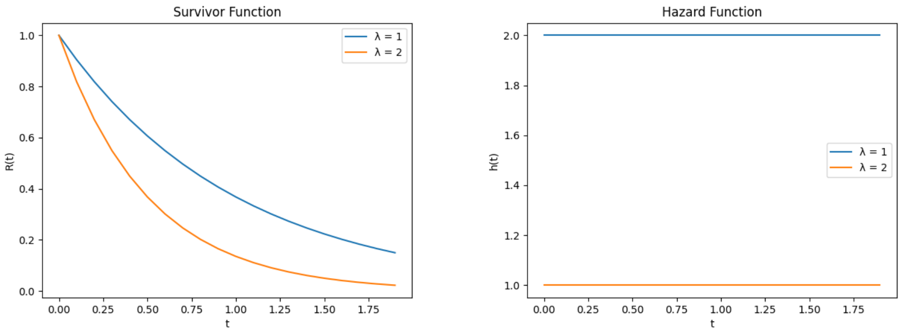

(a) Survivor function for exponential distribution. (b) Hazard function for exponential distribution.

Figure 3.4: Exponential distribution survivor and hazard functions.

This distribution has been widely used to model electronic components [93, 135]. It is based around the assumption that a used component that has not failed is statistically as good as a new component.

Weibull Distribution

The Weibull distribution is a generalization of the exponential distribution. It is well suited to model components that suffer wear-out, such as most mechanical components [82, 135].

$$S(t) = e^{-(\lambda t)^{\kappa}} \tag{3.10}$$

$$h(t) = \kappa \lambda^{\kappa} t^{\kappa - 1} \tag{3.11}$$

Where λ is a positive scale parameter, and κ a positive shape parameter.

The Weibull distribution is adequate to model lifetimes having a constant ($\kappa = 1$), strictly decreasing from infinity ($0 < \kappa < 1$) or steadily increasing ($\kappa > 1$) failure rates.

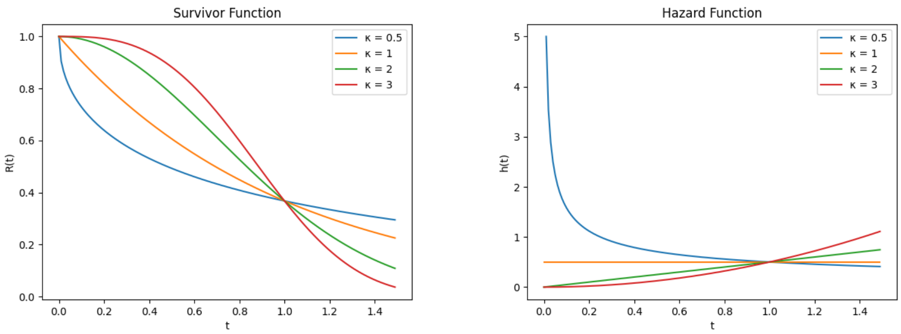

(a) Survivor function for Weibull distribution. (b) Hazard function for Weibull distribution.

Figure 3.5: Weibull distribution survivor and hazard functions.

Parameter Estimation

There are several ways to estimate the parameters of a probability function based on a set of empirical data. Of these, three are highlighted:

- Least Squares Estimator;
- Maximum Likelihood Estimator;
- Bayesan Inference.

The Least Squares Estimator (LSE) is applied by minimizing the sum of the squares of the residuals (SS_{res}) between the parametric reliability curve and the Kaplan-Meier Estimate for a varying parameter vector θ=[θ_1, θ_2, ..., θ_k]^T.

$$SS_{res} = \sum_{i=1}^{n} (\hat{S}(t_i) - R(t_i|\theta))^2 \tag{3.12}$$

$$\theta_{LSE} = \text{arg min } SS_{res}(\theta|t) \tag{3.13}$$

The Maximum Likelihood Estimator (MLE) works by calculating the likelihood L of distribution parameters for the observed events. This is the product of probability densities of n failures at observed times t_i, and the probabilities of survival of m censored items at censoring times t_j, for a given distribution. The parameters that maximize the likelihood are selected through iteration.

$$L(\theta|t) = \prod_{i=1}^{n} f(t_i|\theta) \cdot \prod_{j=1}^{m} R(t_j|\theta) \tag{3.14}$$

$$\theta_{MLE} = \text{arg max } L(\theta|t) \tag{3.15}$$

The Bayesian Inference method is based on the Bayes theorem. A prior conception about the system, under the form of a probability distribution of the model parameters $P(\theta)$ is used to calculate the posterior distribution of these parameters.

$$P(\theta|t) = \frac{L(t|\theta) \cdot P(\theta)}{\int L(t|\theta) \cdot P(\theta) d\theta} \propto L(t|\theta) \cdot P(\theta) \tag{3.16}$$

The Maximum-a-posteriori (MAP) parameters can be calculated:

$$\theta_{MAP} = \textbf{arg max } P(\theta|t) \tag{3.17}$$

Selection of the Estimator

The Least Squares Estimator is a powerful tool, due to its simplicity, robustness and ease of use. However, it is not as appropriate as the MLE or Bayesian Inference methods when considering a more limited set of data and/or a high proportion of censored data [82].

The Bayesian Inference is by far the most complex of the three. Although it is particularly appropriate for limited observations, as it is based on a prior guess, it is very dependent on how well-informed this prior is, it is computationally intensive, and its computation can lead to intractable integrals or an improper posterior.

Given the low quantity of failures in the Falcon 9 data, the LSE method is discarded for this application, and given the lack of informed prior beliefs for the cases considered later in this research, the Bayesian Inference method is also discarded. The MLE is considered therefore the best and easier option in order to have a quick and data-based parametrization of the chosen distributions. In the future, in case vaster and richer data is obtained, the LSE methodology might be appropriate, and likewise for the Bayesian one, in case informed prior beliefs can be derived in the future.

In chapter 5, section 5.4, the MLE method is used to obtain the parameters for the lifetime distributions of the survival function of the Merlin engine, and for the success of the landing operation of the Falcon 9 vehicle.

3.4. Reliability Analysis

Reliability analysis techniques are used to obtain the total system reliability from the estimates.

In this methodology, a mix of reliability block diagrams and fault tree analysis is used in order to obtain a reliability figure for the mission success, from the subsystem level reliability estimates.

3.4.1. Reliability Block Diagram

A Reliability Block Diagram (RBD) illustrates the logic connections between physical components of a system, which makes it possible to translate them into a mathematical model. It allows for trade-offs related to system safety [112].

Series blocks are used when all components are necessary for successful system operation, while parallel blocks are used when only one element needs to operate successfully. The entire system will operate successfully when an uninterrupted path exists between the input and output of the system.

Series system:

$$R_s = \prod_{i}^{n} R_i \tag{3.18}$$

Parallel system:

$$R_s = 1 - \prod_{i}^{n} (1 - R_i) \tag{3.19}$$

Also in parallel systems, the fact that not all components are active at the same time can be accounted for. In this case, some different configurations exist, that depend on the failure rate as well as the reliability of the switch [112].

In an 'm' out of 'n' system, the system operates successfully if a minimum number 'm' of the components is working. A pertinent example to the current study is the "engine-out" operation of a launcher. This will be further explored in section 3.5.4. A complex block diagram relation is used [132]:

$$R_s = 1 - \prod_{i}^{n} \binom{n}{i} R_i (1 - R_i)^{n-i} \tag{3.20}$$

Some disadvantages of RBDs are that breaking complex systems into a block schematic takes considerable effort [112] and implies a substantial number of calculations [132], and that they don't account for situations in which the failure rate changes during the operating time [132], or for non-hardware failures [112].

3.4.2. Fault Tree Analysis

Fault tree analysis is a top-down, deductive reliability analysis method that seeks to determine the probability of a top event happening by graphically representing the sequence of events that leads to an undesirable event [112]. A top event is defined as an event that results in mission or vehicle loss. At the bottom of the tree are the components that make up the system. The result of the analysis is a symbolic representation of the logic behind the events that lead to a mission loss case. It is a method similar to the RBD, with the difference that it calculates the probability of failure rather than the reliability, therefore, it suffers from the same disadvantages with regards to complex systems and dynamic analysis [132]. An advantage of FTA analysis is its risk assessment nature, which also allows for the calculation of the probability of combined faults and failures, and the identification of common cause failures and single point failures. However, they only account for one top event, which has to be predicted and isolated by the analyst, as well as its significant contributors. Furthermore, failure rates have to be available.

In chapter 5, a Fault-Tree is developed for the Falcon 9 vehicle, in order to extrapolate system and launcher level reliability, from the reliability estimates at subsystem and equipment level.

3.5. Reliability Increase

As this study concerns also the trade-off that leads to an increase in reliability, several reliability increase methods are here investigated, from the passive increase in reliability coming from testing and operational experience, to active measures that increase the reliability of a system.

3.5.1. Reliability Growth

Duane [100] was the first to observe and mathematically describe the increase in reliability with cumulative operating time of a system. Reliability increase is noted in both testing and normal operation.

This is due to design changes and modifications throughout both testing and the operating life of the system.

Duane noted in 1964 that the number of failures per total test time, when plotted on a log-log paper, described a linear relationship. The slope of the line is the reliability growth slope.

This has been used by Krevor [132] following a philosophy closer to the Army Materiel Systems Analysis Activity (AMSAA) by Crow model [180], to estimate reliability growth with each launch, where K is a constant denoting the failure rate in the first flight (T=1).

$$h = KT^{-\alpha} \tag{3.21}$$

The reliability growth factor α is taken from recommendations in literature and historical values. A value of 0.2 is used when corrective action is taken for important failure modes, and a value of 0.4 or greater is recommended when there is a program dedicated to failure elimination [132]. Applying the model to known launchers, a result of $\alpha = 0.2006$ was obtained for Atlas, $\alpha = 0.0669$ for Delta and $\alpha = 0.0570$ for Titan [129, 132].

This reliability growth model is useful to apply to reliability increase in development, as well as the increase in reliability that can be anticipated from flying the reusable system. It also addresses the recommendations given by Martino [149] regarding complementing the parametric reliability model with the impact of the qualification program of the engines. How this translates to cost will be explored in chapter 4, section 4.3.2.

Several other more sophisticated reliability growth models exist [92, 94, 111, 116, 131, 188]. The Crow-AMSAA is selected due to its simplicity, applicability, and previous work in the area of launcher and liquid engine reliability where the model was used successfully [132, 170, 196]. It shall be applied in chapter 5 to the Falcon 9 vehicle, in order to predict the reliability increase with the maturity of the system, and in chapter 6, to simulate the reliability increase of a liquid rocket engine in its qualification process.

3.5.2. Testing

The Duane and Crow models for reliability growth can be applied in the testing phase in order to get an estimate for the reliability growth in testing. In addition to that, Pempie [147, 153] suggests a "counting method" in order to relate the reliability achieved with the testing plan using a binomial distribution. This allows to obtain a reliability estimate without the need to have any prior knowledge or assumption about the system, other than the experiments are independent.

$$R = (1 - CI)^{1/MEQ} \tag{3.22}$$

In this equation, CI is the confidence interval and MEQ is the Mission Equivalent parameter, which depends on the testing.

$$MEQ = \omega_1 \frac{\sum \text{cycles tested}}{\text{Mission ignition quantity}} + \omega_2 \frac{\text{Cumulative test duration}}{\text{Mission duration}} \tag{3.23}$$

The ω_1 and ω_2 parameters are weighting factors of the cycle and mission time influence over failure, and can be achieved with expert judgment and Failure mode, effects, and criticality analysis (FMECA) studies. For a first estimate, an equal weighting where $\omega_1 = \omega_2 = 0.5$ is recommended [147, 153].

3.5.3. Redundancy

The increase in reliability resulting from the implementation of redundancy can be easily obtained using the fault tree analysis. 'Hot' redundancy, where two identical systems are operating simultaneously, and both need to fail in order to result in the failure of the element at the higher indenture level, can be easily modeled using a simple 'and' logic gate, where the probability of failure is obtained from:

$$Q = 1 - \prod_{i}^{n} (1 - R_i) \tag{3.24}$$

Whereas the probability of failure of a simple component is modeled at the same indenture level with an 'or' gate:

$$Q = 1 - \prod_{i}^{n} R_i \tag{3.25}$$

3.5.4. Engine-Out

Engine-out capability consists in the operation of the launch vehicle when one engine fails in a benign manner [121, 132]. Paired with a health monitoring system (HMS), this allows the on-board computer to shutdown an engine in case of failure, therefore converting what would be a catastrophic failure into a benign one.

This can be modeled with a conditional 'or' gate, where 2 out of n engines would need to fail in a benign way in order to have this translate into a catastrophic failure. Incorporating the catastrophic failure percentage which isn't avoidable by the HMS, it is possible to model the reliability of an engine cluster with engine-out technology with the following equation, adapted from Krevor and Huang [121, 132]:

$$R_{cluster} = (R_{eng}^{CF})^n [nR_{eng}^{(1-CF)(n-1)}(1 - R_{eng}^{1-CF}) + R_{eng}^{n(1-CF)}] \tag{3.26}$$

On top of this, previous studies [121, 132] suggest including the effect of common-cause failures (CCF) in the reliability model of the engine equipment. Common cause failures stem from systematic defects present in components taken from the same batch of produced units. It is important to account for this, as this type of failures bypass redundancy. This is modeled with a common cause failure factor, represented by β.

$$R_{cluster} = R_{eng}^{\beta} (R_{eng}^{CF})^n [nR_{eng}^{(1-CF-\beta)(n-1)}(1 - R_{eng}^{1-CF-\beta}) + R_{eng}^{n(1-CF-\beta)}] \tag{3.27}$$

In figure 3.6, the advantage of the engine-out solution at reliability level is exemplified. Considering an engine with a reliability of 0.995. As the number of engines in an engine cluster increase, the reliability of the propulsion system goes down. This is modeled with an 'or' gate, as a failure in one of the engines is enough to result in the failure of the whole system. By implementing an engine-out solution, as a 2 out of n gate, the intensity (or slope) of the reliability decrease experienced with the increasing number of engines is effectively mitigated, depending on the catastrophic failure percentage.

3.5.5. Derating

Derating consists in operating an item at a stress lower than its rated design value. Reliability can therefore be increased. Operational derating requires then a design uprating, meaning that the item is designed to support higher stress than that under which it will be operated at. This is an important requirement to not only increase reliability, but also to enable engine-out capability.

From Kim [125], the operational hazard in the derated design can be obtained with the equation:

$$h_{derated}(t) = \mathrm{K}_1 \mathrm{K}_2 h(t) \tag{3.28}$$

Where K_1 is a constant that relates the uprated design failure rate to the non-derated design, and K_2 is a constant that relates the uprated design failure rate to its operation at a lower stress than what is rated for. These constants follow from:

$$\mathrm{K}_1 = \frac{1}{\eta^w} \tag{3.29}$$

$$\mathrm{K}_2 = (1 - \phi) + \phi e^{-\chi(1-\eta)} \tag{3.30}$$

Where $1 - \eta$ is the derating level, w is the exponent of the design uprating for the nominal failure rate, ϕ is the fraction of failures affected by operational derating and χ is a reliability factor dependent on propellant type.

When considering liquid rocket engines from advanced aerospace countries, it is suggested using $w = 0.1017$ [106]. The factors ϕ and χ are taken at 0.2 and 5.78 for kerosene fueled engines, and 0.35 and 12.06 for hydrogen fueled engines [191].

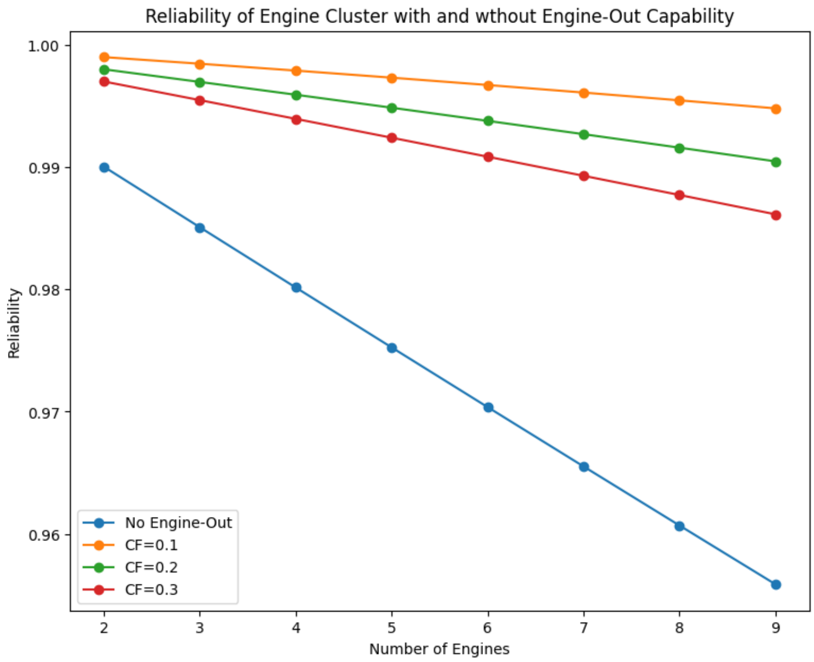

Figure 3.6: Effect of the engine-out capability, depending on the catastrophic failure percentage.

In chapter 6, the effect of derated engine design is explored in more detail, in an analysis comparing a single-engine first stage solution versus multi-engine configurations with engine-out capability and derated design.

3.6. Summary and Conclusion

Reliability modeling is a central part of this study, as it provides the probabilities of failure necessary to assess the expected cost of failure and value of a mission, once combined with the cost estimates.

Modeling shall be done at subsystem level, with special attention given to the propulsion subsystem and in particular to the engine as a component, as they are statistically the most common causes of failure in launches.

Non-parametric and parametric methods were highlighted. For both, access to vast and varied data is important. However, given the nature of the private commercial companies operating the current reusable launch vehicles in the market, it is difficult to lead a purely data-driven study. This research shall therefore hold this as a main obstacle, which in the future can be overcome, therefore refining the results obtained with the same methodology. For now, historical data, publicly available operational information, and expert judgment, shall play a central role in the assessment of reliability. Kaplan-Meier Estimators are recommended, as they are free from assumptions and provide an accurate reflection of the survivor function of a component. They shall be used to also provide a goodness-of-fit comparison with the parametric models, based on lifetime distributions. Exponential and Weibull distributions are highlighted, as they allow to model a great variety of situations and have been used previously in similar works. The maximum likelihood estimator method is selected as the most adequate way to estimate the ruling parameters of these distributions for the purposes of this study.

Reliability analysis techniques will allow to estimate the reliability at the upper levels of the reliability breakdown structure. A hybrid of Reliability Block Diagrams and Fault-Tree Analysis is chosen, based on the previous works on the topic, and the focus of this work on failure.

Finally, attention is given to the reliability increase techniques applicable to launch vehicles. Reliability growth, testing, redundancy and derating design are explored. An engine-out model is adapted and plays a central role in the overall model.

4 Cost Model

The cost modeling approach in this thesis follows the molds of established works at TU Delft, namely SOLSTICE [99]. This model consists of a hybrid approach, mixing a theoretical first unit equivalents method to estimate development and manufacturing costs at equipment level, with CERs from Koelle's TRANSCOST model [129], so as to obtain an estimate for the operations.

In order to expand this model, equipping it to be able to also estimate the cost per flight of reusable launch vehicles, three alterations are implemented:

- New T1 CERs for reusable hardware are used, in order to derive development and manufacturing costs with the the T1 equivalents method. These CERs come from NASA's Crewed and Space Transportation Systems Cost Model (CASTS) [145], and are calibrated using adjustment factors based on data from past ESA missions.
- The recovery costs are included in the direct operations costs. Data from a German Aerospace Centre bottom-up study on recovery of reusable launch vehicles is used to derive CERs for return to launch site, down range landing at sea and in-air capture [172].
- The refurbishment costs model used in previous work [165], combining the effect of increasing cost due to aging with an overall learning effect [165, 184], and applying it to a T1-based CER from the NASA's Operations Cost Model (OCM) [103].

Furthermore, an innovative approach is taken by considering the whole life-cycle of the launch vehicle. This helps better capture the economic implications of the operation of reusable launch vehicles. This is expressed through the following:

- Analysis is done on the whole predicted life-cycle of the launch vehicle, taking into account the effects of aging and learning;
- A new factor focused on the cost of failure is implemented, combining cost modeling with reliability modeling;

The foundation based on SOLSTICE [99] and First Stage Reusability Tool [165] is therefore complemented with CERs coming from the NASA Crewed and Space Transportation Systems Cost Model (CASTS) [145], and the NASA Operations Cost Model (OCM) [103], calibrated with ESA data and SPICE results [161].

The standard currency used in this work is Euro from the fiscal year 2021. Inflation adjustments to data from different years is done with the OECD Consumer Price Index (CPI) values listed in Appendix A, and currency conversion between Euro and Dollar follows the references from the European Central Bank, listed in Appendix B.

Regarding the accuracy of the cost results, internal studies at the Cost Engineering Section at ESTEC showed that in general, the tools and models available to estimate the cost of spacecraft related programs was within the $\pm$20% of the cost of the final contract [97]. This leads to the implementation of a risk factor of 20% on the final estimates for safety. Regarding launch vehicles, as there is less data and more opacity, the risk is usually considered to be 30% on the final contract [152]. While this margin is applied to the final cost, throughout the estimating process, efforts will be made to certify that the items in the cost breakdown also stay under 30% of the publicly available references. Although the errors propagate along the estimate into higher levels of indenture, minimizing the errors along the way ensures that the final estimate will be under the acceptable range of uncertainty.

4.1. Cost Breakdown Structure

Obtaining or building a Work Breakdown Structure (WBS) is one of the first tasks involved in developing a cost estimate [142, 160]. The main objective is to provide a consistent structure that includes all elements of the project that the cost estimate will cover.

The Cost Breakdown Structure (CBS) which will serve as a basis for this study is summarized in table 4.1.

Hardware	Unit	Equipment	Part
	Stage		
		Solid Casing	
		Pres. Tank	
		Fuel Tank	
		Oxidizer Tank	
		Stage Structure	
			Thrust Cone
			Skirt
			Thermal Control
		Engine(s)	
		TVC	
		Press. System	
		Pipes & Valves	
			Pipes
			Valves
		Stage Harness	
	Payload		
		Payload Adapter	
		Payload Fairing	
	Avionics		
		Comms	
		Power	
		Data Handling	
		GNC	
		Avionics Harness	
	Attitude Control		
		ACM	
	Recovery		
		TPS	
		Grid Fins	
		Parachute	
		Landing Gear	
	I&T		
		Stage I&T	
		PAA I&T	
Operations			
	Direct Operations Cost		
		Ground Operations	

		Propellant Flight and Mission Operations Recovery Transportation Fees and Insurance
	Indirect Operations Cost	
	Refurbishment	
Expected Cost of Failure		
	Failure Cost	
		Vehicle and Fight Replacement Insurance Penalty Investigation Implementation of Changes Downtime

Table 4.1: Cost Breakdown Structure.

The hardware related to the typical expendable launch vehicle configuration follows directly from SOLSTICE [99]. The RLV hardware, consisting of Thermal Protection System (TPS), grid fins, landing gear and parachute, are added in order to allow the estimation of these components which are observed in reusable launch vehicles being operated in the present. The operations elements follow from the guidelines in TRANSCOST [129]. Finally, the cost of failure elements relate to the most influential factors in the cost incurred with failure [129, 150].

4.2. Theoretical First Unit Cost

The theoretical first unit cost or T1 estimate is the backbone of this cost modeling strategy. From it, it is possible to obtain the development estimate by using a model equivalence philosophy [99, 138, 161], and the manufacturing costs by applying learning curve theory [88, 99].

The T1 estimate itself is derived from available data, assuming a power law relationship between equipment cost C and equipment mass M [99, 129, 160, 161].

$$C = a \cdot M^b \tag{4.1}$$

The a and b values are linear regression coefficients, obtained from normalized historic mass-cost data points:

$$log(C) = log(a) + b \cdot log(M) \tag{4.2}$$

The coefficients for the CERs used in this research can be found in the relevant documentation of their respective models: SOLSTICE [99] and CASTS [145].

4.2.1. Reusable Hardware Adjustment Factors

The reusable hardware T1 was costed using the CASTS CERs, as the SOLSTICE model does not contemplate items pertaining to the recovery of the launch vehicle stages.

These CERs are based on historic NASA mission hardware, and are obtained in a similar way to SOLSTICE. The main limitation of these equations is that the underlying data points are opaque, on one hand, and on the other, some CERs encompass very different and sometimes inapplicable systems, relying therefore on adjustment factors (AF), to calibrate the estimate to be closer to a particular historical system. This consists therefore in a method which combines the parametric estimation of cost, based on system characteristics, with an analogy method, which brings the result closer to the most comparable historical system. These adjustment factors, when multiplied by the exact CER result, bring the final result of the CER to the corresponding exact cost, therefore obtaining the original data point.

Using ESA data from past and current missions, a new set of adjustment factors were derived using the same philosophy, in order to calibrate the CERs to the closest possible technologies for which there are ESA references, and to the fiscal year of 2021. These factors are summarized in table 4.2.

Element	CER	AF
TPS	Thermal Protection System	3.5
Landing Legs	Recovery Systems	0.2
Parachute	Recovery Systems	0.6
Grid Fin	Mechanisms-Other	0.0013135

Table 4.2: Adjustment Factors developed for use with CASTS CERs, based on ESA data.

Verification

The TPS adjustment factor is in the same order of magnitude of the adjustment factors recommended for systems such as the Apollo Command and Service Module (CSM), Saturn V S-II stage and Space Shuttle Orbiter heat shields. These can be consulted in the CASTS documentation [145]. When applying the estimation method described in [155, 165] to the internal ESA reference, the error obtained was around 91%, and an order of magnitude below the actual result. This result is expected, as the underlying method is a rough generalization, coming from a cost estimate of a specific polyamide felt [157]. The adjustment factor process allows for a better tuning of the CER result in this case.

The landing legs and parachute estimates are derived from the same CER, using different adjustment factors. This is a particularity of the regression done in CASTS, which uses only two data points for the landing gear, coming from the Apollo Lunar Module, and the Shuttle landing gear. The references for the parachutes are the Apollo CSM and the Shuttle SRB.

The adjustment factor obtained for the parachutes is in the same order of magnitude as the references used, however it is slightly larger. This could be due to the higher complexity needed to recover a whole stage, when compared to a capsule or a solid rocket booster. A bottom-up approach such as the one used in previous works [155, 165] is useful when there is extensive knowledge of the system. For instance, when estimating the cost of the SRB parachute system, the error obtained with the adjustment factor here recommended is about 6 times higher than the one obtained with the bottom-up approach. However, for the purposes of this study, information such as the number and length of the suspension and riser lines is not something that is assumed the designer knows at this early phase in the study [136]. This shows however that the engineering build-up method is more accurate, as expected, but more costly from a scheduling point of view.

The adjustment factor for the landing legs is on the contrary slightly lower than the references. This could be due to the lesser complexity of the system when compared to the shuttle gear, on one hand, and on the other hand, the relatively inflated costs associated to the Apollo program (just to the lunar module, a cost of about 23 billion dollars FY2020 is attributed [95]). Given the scarcity of underlying data used to build this CER, the estimate will be as good as the quality of the adjustment factor used. The one provided in this thesis is based on similar systems using ESA data. Employment of expert judgment is particular important for this case, and it is recommended to look at the adjustment factor suggested as an analogy with close historical systems. This methodology allows for a more granular cost analysis when compared to past works [165], where the landing legs cost is folded into the stage level cost estimate.

As for the grid fin estimation, a versatile mechanisms CER is used as foundation, adequate for estimating of miscellaneous mechanical systems. The underlying data points are very diverse, including for instance payload bay doors and docking adapters. The most relevant system for comparison however is the Space Shuttle flight control system actuators like the speed break and the rudder. The adjustment factor recommended here is derived from similar hydraulic systems for which ESA has mass/cost data. It is however two orders of magnitude smaller than the one recommended for the Shuttle actuators. Despite this, the result is not unexpected, as the Falcon 9 grid fins are relatively simpler structures when compared to the Shuttle control surfaces. On top of that, the Space Shuttle Orbiter is a manned system, which involves much higher costs due to its intensive design and extensive qualification to meet the most stringent requirements associated with human space flight [129]. Once again, this result will be as good as the closeness of the system that is being costed to the one from which a comparison through analogy is being made, in order to derive an applicable adjustment factor.

However, it allows in this case for a more detailed analysis, compared to previous work where the grid fins were not part of the cost breakdown, and were folded into the structural cost of the stage [165].

In chapter 5, the recovery hardware of Falcon 9 is estimated using the CERs from CASTS, calibrated with ESA data, and an error of 2% is obtained, when compared to the methodology suggested by Rozemeijer [165]. This method consists of estimating the cost of the landing hardware as part of the structural mass of the stage, using a liquid rocket engine tank CER from TRANSCOST [129].

4.3. Development Cost

In this chapter, the method presented in SOLSTICE [99] that allows to estimate development cost using T1 equivalent costs is described. This is important to revisit, as the failure cost model draws elements from this methodology. On top of that, an equation that relates the cost of a testing and qualification program of engines to its length is developed and discussed.

4.3.1. T1 Equivalence Method

The development cost estimate is obtained through the T1 equivalents method by applying a model development philosophy, and extrapolating its costs in relation to the T1.

This is achieved by determining the different quantity and type of models to be used in the development process, and how they relate to the first unit costs.

The process of estimating the development costs as detailed in SOLSTICE [99] will here be explained in a summarized and top-down manner:

Starting from the Theoretical First Unit cost, or T1, the first Flight Model cost (FM1) can be obtained by removing the management and product assurance percentage ($M/PA\%$). The result is therefore the hardware exclusive fraction of the T1.

$$FM1 = T1(1 - M/PA\%) \tag{4.3}$$

The development cost (DEV) is decomposed in Project Office costs (PO) and Manufacture, Assembly, Integration and Testing costs ($MAIT$), reduced by a chosen cost reduction factor (c_p).

$$DEV = c_p(PO + MAIT) \tag{4.4}$$

The Project Office costs are related to the desk-focused activities, and can be further broken down into its component Engineering (ENG) and Management and Product Assurance (M/PA) functions.

$$PO = ENG + M/PA \tag{4.5}$$

The engineering function of the project office costs is given by the Design and Development (DD) T1 equivalent.

$$ENG = FM1 \cdot DD \tag{4.6}$$

It is suggested in SOLSTICE that the DD factor should be equal to 3 plus the delta TRL involved in the development effort.

$$DD = 3 + \Delta TRL \tag{4.7}$$

The Management and Product Assurance function can be obtained by calculating the corresponding fraction present in the engineering function and in the Manufacture, Assembly, Integration and Testing function:

$$M/PA = M/PA\%(ENG + MAIT) \tag{4.8}$$

The MAIT costs encapsulate the hardware and test routines happening during the development phase. It is a function of the System Test Hardware (STH) costs. It also depends on the number of hardware items $\#HW$ and a potential development cost improvement factor or learning curve factor L_d.

$$MAIT = FM1(STH + L_d \cdot \#HW) \tag{4.9}$$

The system test hardware factor follows from the model philosophy chosen. It is obtained from the number of models utilized, and how they relate to the first unit cost. In SOLSTICE, a Proto-Flight Model approach is selected and recomended, in order to reduce costs . This approach involves a Development Model (DM) corresponding to 0.1 of T1, an Engineering Model (EM) corresponding to 1.3 of T1, and a Proto-Flight Model (PFM), which incorporates the Qualification Model (QM) and some refurbishment to reuse as flight model, amounting to 1.5 of T1. This results in a STH of 3.1.

The development cost build-up based on the flight model cost is best illustrated in the flowchart in figure 4.1.

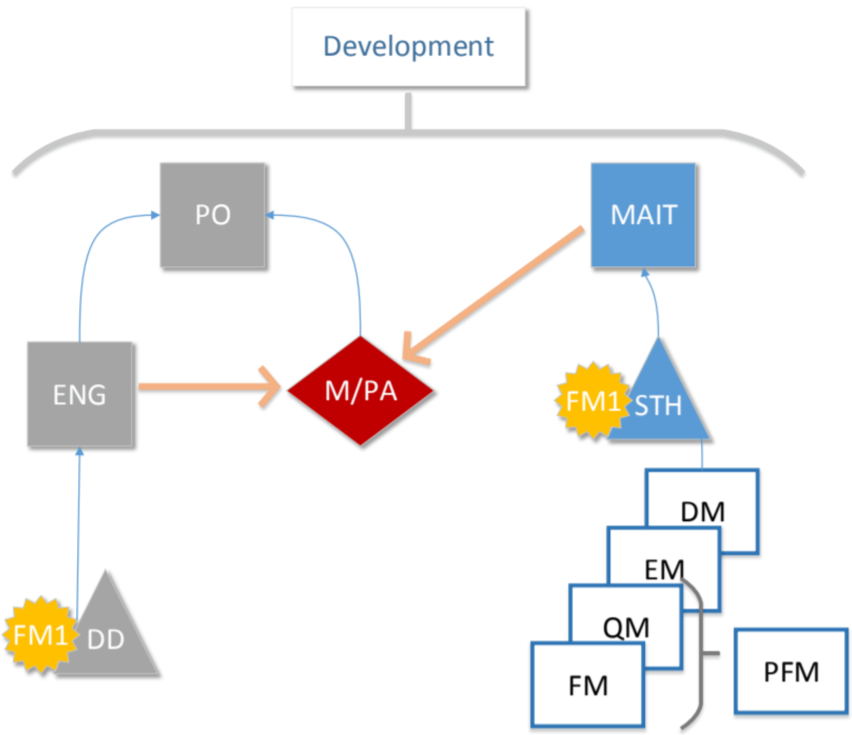

Figure 4.1: Development Cost Build-Up based on First Flight Model [99].

4.3.2. Testing Cost

The cost of testing is implicit in the SOLSTICE methodology, where the hardware costs are estimated through the system test hardware equivalent (STH), but also taking into account the integration and test effort (I&T) as a percentage.

However, this relates to the T1, being therefore a product of historical data and the typical testing effort for comparable systems. It does not take into consideration the possible variations and particularities of the testing program itself, such as extended testing duration and specific additional targets.

From European Space Program data supplied by ESA, pertaining to engine testing duration and costs, a CER was derived, in the form $y = ax + b$, relating the cost of testing to the time of testing, in months. The data included three data points, with costs spanning two different orders of magnitude. The equation derived from three data points is given below (SE=0.3899).

$$C_{testing} = 2.9528 \cdot t_{months} - 4.7816 \tag{4.10}$$

This equation can be used, in conjunction with further information such as the average number of testing per month and total number of tests to achieve a certain target reliability, in order to more accurately estimate the cost of testing, or to estimate an additional testing effort in order to reach a reliability target.

In terms of applicability, this equation pertains to the engine qualification programs using the ESA test stands, such as the P4 and P5 testing facilities of DLR in Lampoldshausen, which have been used to qualify engines of heavy-lift vehicles. It is therefore suitable to be applied to comparable systems. Furthermore, this equation provides negative results for shorter qualification programs, which is not accurate to reality. This isn't unexpected, as engines of heavy-lift vehicles typically have long qualification processes. In order to avoid this problem, when costing shorter test programs, it is suggested considering an engineering-build-up process, taking into account headcount. Alternatively, an option would be deriving a minimum operating cost for short qualification programs.

4.4. Manufacturing Cost

The manufacturing cost estimate is obtained in this methodology by applying a Crawford learning curve [77] to the T1 cost. The cost of the n^{th} unit, MAN_n, is obtained by the relation:

$$MAN_n = T1 \cdot n^b \tag{4.11}$$

Where n designates the unit number and b is the learning exponent, obtained from the learning factor p, the factorial cost associated to production doubling:

$$b = \frac{log(p)}{log(2)} \tag{4.12}$$

A learning factor of 90% is suggested to be used in a first study of launch vehicle systems [96, 97].

An integration and testing factor is also included, ranging between 7.6% and 9% of the hardware cost [99].

4.5. Operations Cost

This section details the operations cost estimating process. TRANSCOST is used for the majority of direct and indirect operating costs, with a greater detail being given to the recovery costs, using CERs derived from a DLR study [172]. Furthermore, different options are explored for the refurbishment cost estimating, including TRANSCOST, empiric models, Rozemeijer's model [165] and NASA Operations Cost Model.

4.5.1. Direct and Indirect Operations Cost

In his "Handbook of Cost Engineering for Space Transportation Systems" [129], Koelle introduces TRANSCOST. This is a top-down model which uses historical project cost data to extrapolate cost estimation relationships and recommendations to estimate the costs related to the whole life cycle and surrounding context of space launch vehicles.

This model was successfully incorporated in the SOLSTICE model by Drenthe [99], in order to obtain an estimate for the direct and indirect operations costs involved in an expendable launch vehicle launch. SOLSTICE relied on public and ESA data to obtain the CERs used to estimate the cost of hardware at equipment level. This therefore goes beyond Koelle in terms of level of detail, which has its hardware scope placed at stage/vehicle level, and only has detailed equations for engine related costs.

However, due to lack of publicly available information and data, estimating the cost of operations at the same level of detail as the hardware estimates was not possible. For that reason, TRANSCOST was used. In this new model for reusable launch vehicles, TRANSCOST will also play a central role in providing a baseline for the estimation of the direct and indirect operations costs. Similarly to previous research, it is recommended that in the future a more comprehensive and focused study on launch vehicle operations and especially the operations pertaining to the modern reusable vehicle configuration is conducted.

The TRANSCOST model often uses the Work-year unit for costs. This is the total company annual budget, excluding subcontracts, divided by the number of productive full-time people. The conversion from Work-year to currency can be obtained using the tables present in Appendix C.

Although refurbishment and spares cost is identified as the third main component of the operations cost, this shall be given special attention in the latter section 4.5.2. The same applies to recovery costs, which while being part of the direct operations costs, will be studied in more detail in section 4.5.3.

The direct operations costs (DOC), as detailed in TRANSCOST, include:

- Ground Operations: This includes the engineering work, site management and support, equipment maintenance (including RLV maintenance), assembly, integration and checkout, launch preparations (tower erection, propellant loading, etc.) and pad refurbishment.
- Propellant cost: Including fuel, oxidizer, pressurants and other consumables, while potentially accounting for boil-off during storage.
- Flight and Mission Operations: This contemplates the costs related to mission planning, evaluation and management, and the launch, flight, tracking and data-relay operations.
- Transport and Recovery: Including the transportation to/from launch site, the recovery and return operations (depending on landing site), and launch assist operations.
- Fees and Insurance: This includes the launch site user fee per launch, the public damage insurance, vehicle loss charge, mission abort charge and other charges.

Ground Operations

The ground operations cost is obtained in WYr parametrically with the following equation:

$$GND = 8M_{GLOW}^{0.67} LpY^{-0.9} N^{0.7} f_v f_c f_4 f_8 f_{11} \tag{4.13}$$

It depends therefore on the vehicle Gross Lift-Off Weight (GLOW) of the vehicle M_{GLOW} in tons, on the launch rate per year LpY,the number of stages N, a vehicle type factor f_v, an assembly and integration factor f_c, learning factor f_4, country productivity factor f_8 and commercial cost reduction factor f_{11}. The values for these factors are summarized in table 4.3

Factor	Condition	Value
f_v	ELV, liquid cryogenic propellant	1.0
	ELV, liquid storable propellant	0.8
	ELV, solid propellant	0.3
	RLV, automated cargo vehicle (SSTO)	0.7
	RLV. Crewed/piloted vehicle (Shuttle)	1.8
f_c	Vertical assembly, checkout on launch pad	1.0
	Vertical assembly and checkout, then transport to launch pad	0.85
	Horizontal assembly and checkout, transport to pad, erection	0.7
f_4	Applied in case of continuous launch operation of $LpY \geq 5$	70%-85%
f_8	Depending on country, 1.0 suggested for first approximation	1.0
f_{11}	Depending on commercialization type, 0.55 suggested by Koelle	0.55

Table 4.3: Ground operations cost factors [129].

Propellant Cost

Given the necessary parameters, the cost of propellants can be trivially obtained with the following expression:

$$PROP = \frac{M_p}{O/F+1} c_{fuel} + \frac{O/F \cdot M_p}{O/F+1} c_{ox} + M_{pres} \cdot c_{pres} \tag{4.14}$$

Where M_p and M_{pres} designate the mass of propellant and pressurant, O/F is the propellant mixture ratio, and c_{fuel}, c_{ox} and c_{pres} represent the specific costs of fuel, oxidizer and pressurant respectively.

This can be further refined by including a parameter for boil-off of cryogenic propellants.

Flight and Mission Operations

The flight and mission operations cost is obtained in WYr parametrically with:

$$FM = 20(\sum Q_N) LpY^{-0.65} f_4 f_8 \tag{4.15}$$

In this equation, Q_N is a specific value that depends on the vehicle complexity:

- Small solid motor stages: $Q = 0.15$ (each);
- Expendable liquid propellant stages or large boosters: $Q = 0.4$ (each);
- Recoverable or fly-back systems: $Q = 1.0$ (each);
- Unmanned Reusable orbital systems: $Q = 2.0$ (each);
- Crewed orbital vehicles: $Q = 3.0$ (each);
- Expendable lunar transfer vehicle: $Q = 2.0$ (each);

Transportation and Recovery Costs

The recovery portion of the direct operations costs is explored in detail in section 4.5.3.

The transportation cost estimation follows from an ESA internal reference used in SOLSTICE, which is 5.837€/kg (2021). This result can be made more precise in a bottom-up estimation, when knowing the type of transport (air, naval, land) and distance of transport.

Fees and Insurance Costs

Three factors are considered here: the launch site user fee, the public damage insurance (third party liability, not to be confused with insuring the launch, the payload or the launch vehicle) and the payload charge fee.

The launch site user fee is suggested in SOLSTICE at 1.327 M€ (2021). This number depends on a number of factors [44], ranging from the location, the types of propellant used, the target orbit, if the vehicle is expendable or reusable, if the launch includes a reentry vehicle, the ground station requirements, etc. These can be further refined in later phases when the characteristics of the launch vehicle are defined.

The public damage insurance is suggested by Koelle to cost 100k€, for a premium coverage of 100M€.

The issue of launcher insurance is further developed in section 4.7.2. According to current data, it is between 4-5% of the cost per flight [5, 41, 197], and is dependent on the launch vehicle reliability [5].

The payload charge fee is suggested in TRANSCOST to be 5.995€/kg (2021).

Indirect Operations Costs

The indirect operations costs are not directly related to the launch. They include staff and administrative personnel costs, marketing and technical support, the so called ”commercialization costs” of the business [129].

In SOLSTICE, a CER for the indirect operations cost in WYr is derived from the graphical interpretation of the results present in TRANSCOST, depending on the yearly flight rate LpY and the percentage of work subcontracted σ:

$$IOC = (33\sigma + 32)LpY^{-0.379} \tag{4.16}$$

4.5.2. Maintenance and Refurbishment

Refurbishment is often used as an umbrella term that includes all the activities done on the reusable hardware in order to have it operational for the next flight. It can include cleaning, repair, maintenance, repackaging, replacement of expended parts, periodic overhauls/heavy maintenance, and acceptance testing.

Koelle [129] makes the distinction between ”refurbishment” and ”maintenance” in the 2013 edition of TRANSCOST. The former is used to describe major overhaul activities that happen offline, when the vehicle is taken out of service after a certain number of flights to undergo detailed inspection and replacement of components before wear-out. Maintenance is defined as the activities that are done online, between two consecutive flights. This definition is echoed in the NASA Operations Cost Model.

In the current panorama, the term ”refurbishment” has been used colloquially to describe the online turnaround maintenance activities done on the Falcon 9 booster between flights.

In this research, the terminology chosen is the following: "refurbishment", "maintenance" and "online refurbishment" shall be used to signify the online activities between flights, included in the vehicle processing flow, as used in the state-of-the-art and especially when referred to the SpaceX vehicles. The term "overhaul" or "offline refurbishment" shall be reserved to refer to the activities that imply removing the vehicle from its normal processing flow, in order to execute a long offline process that can include inspection, repairs, spare replacement, overhauls and acceptance testing.

Online Refurbishment

The online maintenance activities are contemplated in the TRANSCOST CER used for ground operations (equation 4.13). These costs replace in part the extensive assembly and checkout activities for expendable launch vehicles. Another method, also proposed in TRANSCOST [129] and implemented in the First Stage Reusability Tool (FRT) [165] by Mark Rozemeijer, is taking the refurbishment effort as a percentage of the first unit production cost. The main challenge with this method is selecting the correct refurbishment cost fraction, as this value can range between 3-50% of the T1 cost [120, 129, 165].

In Rozemeijer's work, a refurbishment fraction of 25% is implemented. On top of that, a "learning curve" effect is also included, that expresses the increase in refurbishment costs that result from the aging of the system. This can range between 105-115% [184].

$$REF_{FRT} = 0.25T1(n^{\frac{ln(1.15)}{ln(2)}}) \tag{4.17}$$

In the NASA Operations Cost Model, the online maintenance is estimated based on a headcount CER included in the launch operations model, which is equivalent to the TRANSCOST ground operations CER.

Offline Refurbishment

The offline refurbishment activities include:

- Detailed inspection;
- Structural hardware replacement;
- Rocket engine replacement;
- Exchange of feed and pressurization system components.

The Space Transportation System is the main example used in TRANSCOST to illustrate the resources required in the offline refurbishment process:

- The orbiter refurbishment was executed every 4.5 years or every 6 flights, lasting for 18 months and 235 jobs.
- The Space Shuttle Main Engine (SSME), due to its complexity, sensitivity, and uprated operation at 107% of the nominal thrust, had to be refurbished (ie, taken out of the vehicle processing flow for offline refurbishment) after every flight. These costs per flight amounted to 11% of its production cost
- It is pointed out that the derated operation of the engine prolongs the life time of the engine and therefore reduces refurbishment costs.

Furthermore, the example of the NK-33 engine is given. This engine was designed as an expendable engine but was qualified for 20 flights on the Kistler K-1 reusable launch vehicle, with refurbishment after 10 flights. An improved version was studied with 75 lifetime flights, with 25 flights between refurbishment.

In 2013, Koelle cites a Rocketdyne study that hypothesizes future RLV engine refurbishment cost per engine would amount to 240 Wh every 20 flights plus 10% spares. With the use of a health monitoring system and data-informed maintenance requirements, this effort is predicted to be less than 0.5% of the first unit production costs per flight, with refurbishment every 20-25 flights, and a lifetime of 40-80 flights.

In the NASA Operations Cost Model, offline refurbishment activities are estimated with parametric CERs based on a first unit equivalent. They include reduction in cost from both learning and rate effects.

Verification

The only system currently being refurbished and reflown is the first stage or booster stage of the Falcon 9 launch vehicle, by SpaceX.

The maximum number of reuses to date has been 12 flights, with booster B1051. This is still far from the expected number of flights before (offline-) refurbishment. There is no record of major overhauls or refurbishment happening (other than engine upgrades), and the turnaround time for booster B1051, ranging between 38 and 231 days, does not indicate major refurbishment either. The minimum turnaround time for a Falcon 9 booster has been 27 days, and the average has been 112.8 days, with a standard deviation of 116.9 days.

The maintenance and refurbishment of Falcon 9 isn't transparent, but some information can be found online, in official communications by SpaceX, and statements by Elon Musk.

Gwynne Shotwell, the President and COO of SpaceX, said in 2017 that the maintenance and refurbishment objectives were to have each booster fly 10 times with no hardware changes, and at least 100 times with moderate refurbishment done within 24 hours [159]. This short turnaround time has not yet been observed.

Hans Königsmann, former Vice President for Mission Assurance for SpaceX, stated in 2018 that most of the stage maintenance work between flights is focused on the engine [8], with the main tasks being part replacement (as preventive maintenance) and tank inspection (which has mostly shown no need for cleaning). Occasionally, TPS has been reinforced.

Further evidence that the Falcon 9 flight-proven boosters aren't removed from the vehicle processing flow back to manufacturing (California) or testing (Texas) is that maintenance operations are done at the launch site in Florida [37, 38].

Elon Musk has claimed that in order to achieve the upper limit of life time flights, the offline refurbishment needed would involve replacement and upgrade of engine parts and cleaning of the engine turbines, which is a difficult process due to the type of propellant and cycle of the Merlin engines. The Methane fed Raptor engines for SpaceX's new generation of launchers aims at solving this problem [29].

Elon Musk has claimed that the costs of ”refurbishing” one booster for its next flight are around $250,000 [23], in May 2020. By this time, 51 boosters had been successfully recovered.

With this information, a comparison of the different models and recommendations is executed:

It is assumed that the Falcon 9 maintenance costs are dominated by the engine maintenance. The engine considered weighs 470 kg, which allows us to provide a T1 estimate of $4.98 M (2020). In order to compare against the figure given by Elon Musk, it is assumed that the ranking on the engine maintenance learning curve is $52 \cdot 9 + 1 = 469$. Starting with the maintenance models based on taking the maintenance/refurbishment costs as a percentage of the first unit production costs.

Maintenance Fraction	Cost Estimate (M€)	Error
3%	0.527	110.8%
10%	1.757	602.8%
20%	3.515	1306%

Table 4.4: Verification of refurbishment models based on percentage of T1.

As can be seen from table 4.4, the errors obtained using the $4.98M T1 estimate for the engine and the assumed ranking on the learning curve of 469 are superior to 100%. Using this methodology and assuming that the refurbishment figure of $250k is accurate, then the refurbishment fraction should actually be 1.4% of the T1, which is a value over 50% lower than the best case of 3% taken from the X-33 experience in [74, 120]. Compared to other vehicles, this percentage of 1.4% is comparable to that registered for an Orbital Scramjet Vehicle (1%), and the Shuttle Orbiter (1.3% for the vehicle, plus 1% when including the engines) [129]. Despite that, it is still a fraction which is two orders of magnitude greater than that of a fighter jet (10^{-2}), and three of an airliner aircraft (10^{-3}) [74]. A Boeing 747 has a refurbishment fraction of 0.006% [129].

Next, the Rocketdyne and TRANSCOST recommendations for offline refurbishment are analyzed. Assuming a quote of $150k for work-year costs [97], Rocketdyne's expectation for reusable launch vehicle refurbishment applied to Falcon would be $153k every 20 flights, which comes down to $7.65k per flight once amortized. With TRANSCOST expectation, the value is $2.241M every 20-25 flights, which amortized is \$89.64k-\$112.05k per flight.

Following that, the OCM models are compared to the \$250k figure. The online vehicle processing refurbishment/maintenance activities estimated at a launch rate of 26 flights per year (the flight rate of SpaceX in 2020) is \$1.059M per flight (323.7% error). The offline refurbishment model's results are summarized in table 4.5, with an operational experience of 84 flights being assumed.

Activity	Cost per Flight (\$M)	Error (%)
Refurbishment (Cleaning and Repackaging)	0.143	42.8
Spares (Repairs and Replacement)	0.021	91.6
Overhauls	0.008	96.8
Acceptance Testing	0.164	34.4

Table 4.5: Verification of offline refurbishment models from NASA OCM.

Given this analysis, it seems like the most applicable relation is the "Refurbishment" CER from NASA OCM, which is taken from the Solid Rocket Motor (SRM) experience and provides a result once divided over each flight with an error of 42.8%. Considering that the stage refurbishment process includes more activities than simply the engine refurbishment itself (namely the stage and potentially fairing refurbishment, line replaceable units, etc), the error should be smaller when these factors are included. This is congruent with the statements reported by the SpaceX officials, which described the Falcon 9 refurbishment as consisting mainly of replacement of parts, with occasional repairs.

If the aging effect proposed by Rozemeijer is applied to this case, the average cost of refurbishment over a cycle of 10 flights centered around the same ranking of the learning curve taking an aging effect of 115% is \$196k per flight, which translates into an error of 21.62%. If the statement that the Falcon 9 booster refurbishment consists mainly of refurbishing the engines, then this model is appropriate, as it lies within the standard 30% for cost estimates. Due to its acceptable accuracy and dependency on the T1 costs, the OCM model will be used in chapter 5 to estimate the refurbshment costs of the Falcon 9 first stage.

In case a model based on a fraction of the T1 production cost is preferred, it is recommended that for a Falcon 9 type vehicle, a percentage of 1.4% be used. Although it is a distant value from what was in 2006 considered the best case scenario [120], coming from the experience of the X-33, a suborbital vehicle, it is considered plausible that a vehicle operated 20 years after the development of the X-33, which was designed for reusability, will have improved the refurbishment cost fraction.

Limitations

Although the NASA OCM provides a satisfactory estimate of both the maintenance and refurbishment costs, with some degree of flexibility, this model has some limitations:

Firstly, when used in conjunction with TRANSCOST, there could be some intersection and double estimating of maintenance costs, since the ground operations CER is the same for both expendable and reusable launch vehicles, justified by the assumption that the cost of maintenance of RLV is compensated by a less costly assembly and checkout process. Using these two models side by side implies the assumption that the ground operations for expendable and launch vehicles are the same from a cost standpoint, and that RLVs have an added maintenance cost.

Secondly, TRANSCOST suggests that with increasing time between refurbishment, its cost should increase. The OCM model does not contemplate this variation, providing simply a yearly estimate. The aging effect mitigates this.

4.5.3. Recovery

Determining the cost of recovery is a big obstacle in the research projects concerning the cost of reusable launch vehicles. This is due to the scarcity of data related to novel landing strategies employed since the Shuttle, such as offshore vertical landing and mid-air retrieval.

At TU Delft, previous works looked to determine to cost of retrieval in a bottom-up fashion:

Pepermans [155] uses data obtained by Snijders [169], in order find an estimate for the cost of retrieval. A constant figure for yearly recovery costs is considered, depending on the recovery strategy.

Snijders data was based on expert advice by employees of Airbus Helicopters for the mid-air retrieval option, and Zwijnenburg Shipbuilding and Damen Shipyards for the ship-based recovery. The

Table 4.6: Retrieval cost model used by Pepermans [155].

Pepermans Retrieval Model	
Type of Landing	Cost per year
Sea	0.760 M€ [169]
Air	0.775 M€ [169]
Land	0.180 M$ [69]

Table 4.7: Retrieval cost model used by Rozemeijer [165].

Rozemeijer Retrieval Model	
Type of Landing	Cost per launch (MY)
Sea	$1.2t_{travel}20/T_{MYrH} + 760000/(LpY \cdot C_{MY})$
Air	$1.2t_{travel}3/T_{MYrH} + 775000/(LpY \cdot C_{MY})$

cost figures obtained include a yearly estimate for buying/rebuilding the recovery vehicle, the labor involved in the recovery, the fuel (diesel for the ship and kerosene for the helicopter), maintenance, and a margin for unforeseen costs.

The estimate for the land retrieval method follows from a high-level estimate for the costs of operating a truck in the US. This is taking into account the average operating cost per mile and a typical figure for the amount of miles covered by a truck on a yearly basis in the US [69].

Rozemeijer [165] further refines this model in its ship and helicopter recovery methods, by adding a factor proportional to the time of travel (t_{travel}), taking into account the the size of the crew (20 for the ship, and 3 for the helicopter) and adding an uncertainty parameter for the duration of operations of 20%. This is then related to the landing range, by taking into account the speed of the recovery vehicles (13.5 knots for the ship and 140 knots for the helicopter).

Although this refinement directly relates to the particularities of each recovery strategy, and provides parametric and bottom-up justifications for the costs of labour, relating it to the distance of the recovery, it fails to subtract from the original yearly figures obtained by Snijders the assumed costs of labor. This results in an inconsistent and redundant accounting of the labor costs.

The Koelle and Rozemeijer models provide a good first approximation for the cost of retrieval. However, they assume *a priori* a fixed yearly retrieval cost, which only depends on the recovery strategy, but not on the properties of the recovered system. This is a main driver of the recovery cost, as in order to optimize cost, the strategy employed to recover a microlauncher would be drastically different in cost as one used to recover a heavy lift stage.

On top of that, these high-level figures are dominated by the cost of the recovery vehicles themselves, and their specific direct operation costs. The indirect operation costs of the recovery vehicles, mission control costs and overhead costs are not included.

At the German Space Agency, or Deutsches Zentrum für Luft- und Raumfahrt (DLR), a more detailed bottom-up study of recovery options for a reusable first stage was performed [172]. This showed that the dominating contributor to the overall cost of recovery, in both the ship and the helicopter solutions, are their indirect operating costs. This might hint at the possibility that the previous models severely underestimate the recovery costs by disregarding these factors and focusing solely on the direct operation costs.

Furthermore, none of the previous TU Delft models consider the return to launch site strategy employed by SpaceX. This method results in a much lower cost of recovery, dominated by the cost of facilities, as shown by the DLR study [172].

The approach recommended by Koelle [129] in TRANSCOST is a good candidate for this study. It is a parametric CER based on the mass of the recovered system (M) and the flight rate (LpY). It was obtained from three reference cases: The shuttle Solid Rocket Booster (SRB) recovery, Ariane I recovery studies performed by sea-recovery company in Hamburg HARMS for ESA/CNES, and the 1963 ROMBUS studies for the recovery of stages/boosters from the sea, by the Douglas Company.

$$C_{rec} = 1.5/LpY \cdot (7L^{0.7} + M^{0.83}) \cdot f8 \cdot f11 \tag{4.18}$$

The biggest limitation of this CER is that it is independent of the recovery method. Furthermore, the reference cases on which it is based are solely sea-recovery methods after a splash-down landing. In addition, the Shuttle experience inserts an upward bias compared to the expected results.

Considering an f8 factor of 1 and an f11 factor of 0.55, the recovery costs as a function of the recovery mass is represented in figure 4.2, for a flight rate LpY of 10, 20 and 30 flights per year.

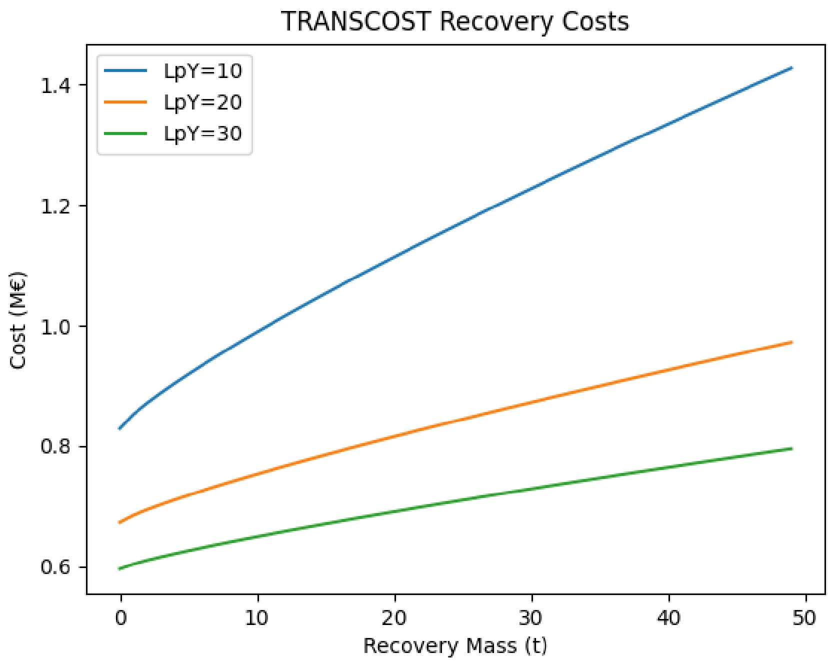

Figure 4.2: TRANSCOST recovery cost model [129].

The DLR team [172] state that in their verification, the upward bias in TRANSCOST was identified, and that vertical landing methods display a higher insensitivity to the landing mass variable, as the resultant increase in cost was negligible. Furthermore, the return to launch site solution was found to be more resistant to changes in the launch rate, as it was estimated at $250k at 15 launches per year, which falls to just below $200k with a launch rate greater than 20 launches per year.

Taking the data points resulting from the DLR study, and updating them to 2021 euro values, three power-law CERs were deduced: SpaceX method for Down-Range Landing at Sea (DRLS), Return to Launch Site (RTLS) and In-Air Capture (IAC):

$$C_{rec_{DRLS}} = 4.713L^{-0.745} \tag{4.19}$$

$$C_{rec_{RTLS}} = 1.518L^{-0.696} \tag{4.20}$$

$$C_{rec_{IAC}} = 10.794L^{-1.036} \tag{4.21}$$

These CERs obtained from the data derived with the DLR method are depicted in figure 4.3. It is visible the conclusion derived by the team that when considering more than 15 flights per year, an in-air capturing method is preferred, provided that the horizontal velocities allow it.

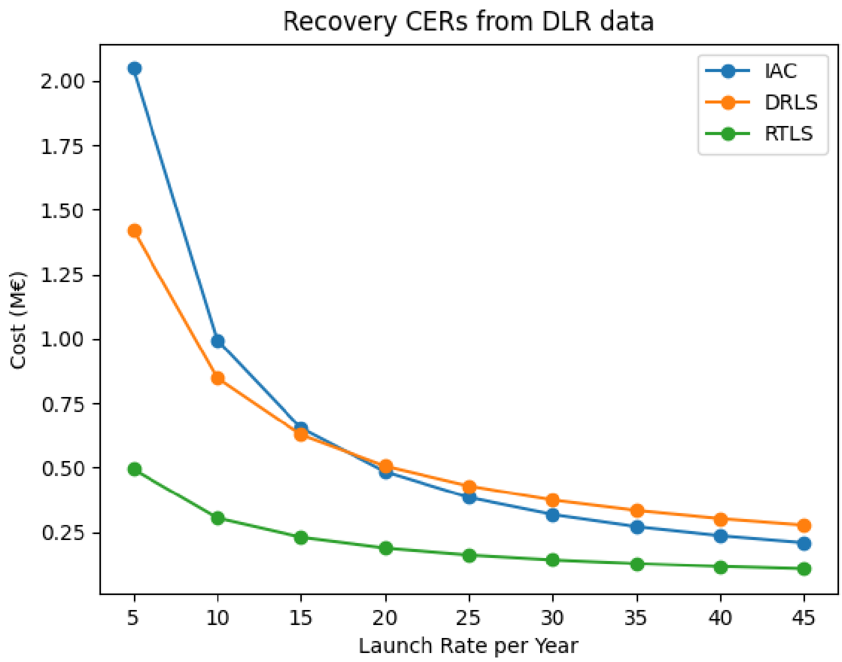

Figure 4.3: Recovery CERs obtained from DLR data [172].

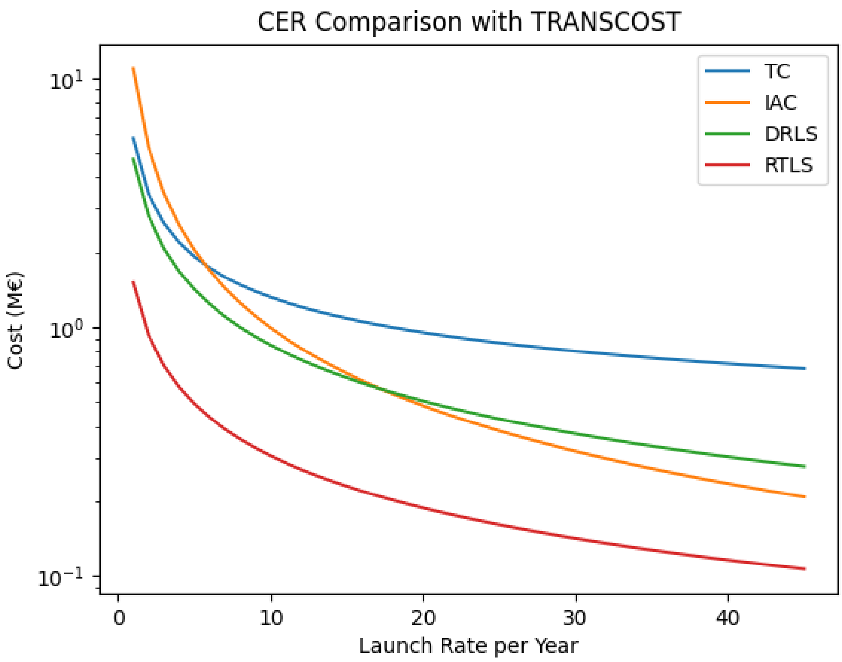

Figure 4.4: Comparison of recovery CERs with TRANSCOST model.

Comparing these results with a TRANSCOST computation, included in figure 4.4 assuming the same recovery mass for a Falcon 9 first stage, it is visible that the costs are in the same order of magnitude. It is therefore considered that the TRANSCOST results are suitable to be applied to a post-splashdown recovery, and that this model may be used for cases where recovery mass is a driving factor in the recovery process. For the remaining cases, where mass of the reusable system is not a driving factor [172] the CERs based on the DLR model shall be used. They were derived using a recovery mass of about 50t, and it was determined that the landing mass on the mission is negligible due to low direct launch costs in all cases. Therefore, their application the Falcon 9 vehicle, which has a recovery dry mass of 25t, should be acceptable.

4.6. Cost and Price per Flight

The Cost per Flight (CpF) is obtained from the conjunction of the three main factors identified: development (DEV), manufacturing (MAN) and operations (OPS).

The development cost is a non-recurring cost, amortized over a certain number of flights. An internal reference at ESA shows a precedent for considering an amortization period of 10 years. Given an annual launch rate of LpY, the development amortization parcel DEV_i pertaining to launch i can be obtained with the following equation:

$$DEV_i = \begin{cases} \frac{DEV}{10LpY} & i < 10LpY \\ 0 & i \geq 10LpY \end{cases} \tag{4.22}$$

The manufacturing cost is a recurring cost, related to actually building the components of the launcher to be flown. It depends on the flight number, due to the accumulating learning curve effects.

The operations cost is a recurring cost. It also depends on the flight number, due not only to the learning accumulated, but also to the aging effect connected to the refurbishment process.

$$CpF_i = DEV_i + MAN_i + OPS_i \tag{4.23}$$

The price per flight is made up of the cost per flight plus a profit margin. It is the final cost presented to the customer. The profit margin is determined at the discretion of the launch provider. At ESA, a cost-based price per flight is instituted, with a profit margin of 8% of the flight cost [175]:

$$PpF_i = 1.08CpF_i \tag{4.24}$$

4.7. Failure Cost

A mission failure results in costs that surpass the simple replacement of parts and possible relaunch of the payload.

This is brought to attention by Koelle in TRANSCOST [129]. Parkinson also puts this topic at the forefront of his paper ”The hidden costs of reliability and failure in launch systems” [150]. Koelle argues that a failure or mission abort can result in costs that are 2-3 times higher than that of a nominal mission, while Parkinson cites a potential factor of 3-5.

In an expendable launch vehicle failure, the launch provider would simply need to supply the client with a free flight to re-launch the payload, in case this was agreed as part of the launch contract. The cost is more severe in the case of a reusable launch vehicle, as the future value generated by flying the reusable components multiple times can be lost.

However, a mission failure implies more costs than simply supplemental manufacturing and operating costs [129, 150]. Standard procedure involves an investigation in order to find the cause of failure, and determine corrective measures. These can result in extra development effort. On top of that, there are costs related to the down-time that results from the failure, as well as other factors, such as climbing insurance rates and loss of business. This is due to the offline period but also to the loss of trust in the system from the customer’s perspective.

The main factors that contribute to this increase in costs include:

1. The costs of replacing the vehicle/flight;
2. For the reusable case, the cost of losing the reusable hardware;
3. Climbing insurance rates;
4. The cost of running an investigation campaign;
5. The cost of technical improvements (development, implementation and validation) and revision of processes;

6. Fixed costs related to maintaining the fleet and its operability (facilities and personnel);
7. Loss of revenues from downtime and loss of business.

The failure cost is obviously an extremely important factor in the calculation of the expected cost of failure. Given its importance to achieving the research objective of determining the cost of failure, a failure cost model is here detailed, which will be validated against the recommendations found in literature.

4.7.1. Vehicle and Flight Replacement

Co-launching parties customarily sign cross-waivers agreeing not to sue each other in the event of the failure [109]. This means that there is no obligation on the part of the launch service provider to offer a re-flight. A re-launch guarantee can however be offered by the launch service provider, for a premium. This will be further explored in the next section.

In the case of a reusable launch vehicle, unless the flight of the failure corresponded to the last flight before decommissioning the vehicle, then there is an added unaccounted cost of building a replacement reusable system.

The replacement cost of the vehicle consists of the cost of manufacturing the reusable system, according to its appropriate rank i on the learning curve (or ranks, when accounting for part redundancy):

$$C_{v.replacement} = MAN_{rs_i} \tag{4.25}$$

In case a re-flight guarantee was agreed to, then there is an added cost of re-flying the payload. The replacement cost of the flight using an expendable booster can be trivially obtained, by calculating the cost of the next flight, taking learning in manufacturing into account. This can be obtained with equation 4.26.

$$C_{f.replacement} = MAN_i + OPS_i \tag{4.26}$$

For immediate purposes, the average manufacturing and operating costs can be used.

$$C_{replacement_{avg}} = MAN_{avg} + OPS_{avg} \tag{4.27}$$

It is important though to mind the fact that if the average cost per flight is used, the amortization of the development cost must be discounted, as it is not value ”lost” with the failure, and so it doesn’t need to be compensated.

Furthermore, it is important to account in the overall simulation that the cost per flight includes a term pertaining to the re-launch guarantee, which is expressed as a percentage of the cost per flight.

4.7.2. Climbing insurance rates

The space activities insurance market is a very extensive and complex topic, which ties into cost, reliability and failure.

In its reports, the FAA distinguishes six types of insurance [109, 110]:

- Pre-launch insurance: covers damages to the spacecraft or launch vehicle during construction, transport and processing;
- Launch insurance: covers losses during launch, from launch failures to improper orbit placement;
- In-orbit policies: insures the satellite during its operational time, against technical problems and damages;
- Third-party liability/government property: protects the launch service provider and the customer in the event of public injury or government property damage. These are included in the model described in 4.5.1.
- Re-launch/Re-flight guarantees (RFG): this is a type of insurance where the launch provider acts as the insurance provider to its customers. In case of a failure, the re-launch or cash-back are assured to the customer in exchange of a fee.

- Business Insurance: covers the loss of revenue in case the satellite fails to reach operational status.

Out of these, the pre-launch insurance and in-orbit insurance are not directly related to the performance of the launch vehicle during flight. The liability insurance as applied to the launch is already covered in the model. The business insurance is interesting, as if the launch fails, the launch service provider will miss out on revenue, similarly to the satellite operator. This is further explored in section 4.7.5. However, in 1998, it was noted that for the satellite case, the premiums for this type of insurance were becoming too expensive and therefore rarely underwritten. In the updated FAA report in 2002, this insurance isn't listed as a common type of space insurance. From an expert judgment consultation, it was found that this is not a common insurance in the space insurance market today [89].

The re-launch guarantee ties directly into the previous section. Since it is customary for co-launching parties to sign cross-waivers agreeing not to sue each other in the event of the failure [109], there is no obligation on the part of the launch service provider to offer a re-flight. In the event of a failure, the launch service provider brunts the cost of manufacturing a new launch vehicle to replace the one destroyed, but has no obligation to re-fly the payload and incur on operating costs. This can be previously agreed with the customer however, in exchange of a premium.

Concerning launch insurance, the satellite operator or the launch service provider have the option of insuring the assets to cover the potential losses coming from of a failure. Insurance premiums are calculated as a percentage (or rate) of the asset insured. In the case of a launch vehicle, insurance has been modeled in relation to the cost per flight [184, 185]. In reality, insurance of expendable launch vehicles is usually limited to liability and pre-flight operations, and the only reusable launch vehicle currently operating, Falcon 9, does not take out insurance policies covering their hardware or potential losses from failure [89]. This can however change in the near future, as demand for new insurance products rises in the space sector [89].

The insurance market has seen rates go down overall over the decades, with some periods of growth. Depending on the launch vehicle insured, in western countries, insurance rates were around 25% in the 80s, 15-17% in the 90s, and in 2020, insurance for vehicles such as Ariane 5 and Falcon 9 were reported to be around 4-5% [5, 41, 197]. This is due to these vehicles high reliability and proven flight record. Current rates for Ariane 5 and Falcon 9 could stand at a mark as low as 2-3% [89].

In the aftermath of the Falcon 9 failure on the 3rd September 2016, it was suggested that insurance rates could double, on one hand as a result of the distrust in the launcher caused by the failure, and on the other as a compensation to the speculation in the launch vehicle insurance market, caused by the good performance at the time of the Ariane 5 and Vega launchers [41, 56]. In this sense, it is expected that this increase is felt across the board in the industry regardless of the vehicle that suffers the failure, which is something that is complicated to model. However, it is expected that in the future insurance, rate increases due to failure may be restricted to the offending vehicle [186].

Good performance and proven reliability lead to a decrease in the insurance rates [110, 129]. It is however hard to model insurance after launch vehicle reliability, as there is a multitude of obstacles. Firstly, the process is long and costly for the insurance provider, given that it depends on so many factors. It also requires a high launch rate per year in order to determine launch vehicle performance, which until recently wasn't seen in the market. Furthermore, opacity in the form of International Traffic in Arms Regulations (ITAR) and commercial interests lead to difficulties in determining reliability [186]. On top of this, the insurance market is reactive to external factors [89], such as the aviation industry, the global economy, war, etc. For that reason, it is very hard to predict and model the value of insurance rates in the future.

The additional costs expressed by climbing insurance rates as a result of failure are propagated in time, as future flights will be subjected to the updated higher insurance rate. These mostly affect directly the satellite operator, as they will face a cost increase to insure their satellite. However, this effect is passed to the launch service provider, as the launch vehicle can become less attractive to the customer if the flight is more expensive to insure (both for the flight insurance and for the re-flight guarantee). If in the future having the launch provider insure the reusable launch vehicle becomes common-practice, then it is expected that a failure will directly result in an increase in costs due to the increasing insurance premiums.

The magnitude of these additional costs depend on the mathematical formulation of the evolution of the insurance rate with time, which isn't known and is very difficult to model, for the reasons that have already been mentioned.

Considering a hypothetical stationary insurance with rate ins, covering the reusable hardware. Assuming that this rate is doubled after a failure, and takes n flights to recover its original value with a linear evolution. The cost due to climbing insurance rates during that period is given by:

$$C_{insurance} = \frac{ins \cdot n}{2} \cdot MAN_{rs_i} \tag{4.28}$$

Furthermore, it is assumed that the given insurance figure for modern launch vehicles is 4% [5, 41, 197], and that it takes 20 flights to recover the original insurance premium. This is the number of Falcon 9 flights since its first commercial launch in 2016, to 2020, when these insurance figures were reported. Given these assumptions, the cost of insurance rate penalty is calculated to be 0.4 of the average cost per flight:

$$C_{insurance} = 0.4 \cdot MAN_{rs_i} \tag{4.29}$$

4.7.3. Failure Investigation

Typically after a critical failure in a launch mission, an inquiry team is established in order to run an investigation into the causes of the failure and the steps to take to correct it. This often presupposes a downtime for the operation of the launcher fleet. In the case of Falcon 9 flights, this procedure is encapsulated in their own FAA approved contingency response plan, and is consistent and compliant with NASA's contingency action plan as agreed per the Commercial Resupply Service contracts [176].

Looking at the investigations for the Ariane 5 vehicle, their duration after failures ranged between 24 and 36 days, with an average of 28 days [15–18]. The two Vega failures resulted in investigations lasting 28 and 55 days respectively [71, 72]. For the Falcon 9 case, the investigations were much longer, with one lasting for 122 days related to their 2016 accident, and one in 2015 lasting for at least 4 months [64, 65]. This is likely due to the fact that losses in reusable vehicles are more costly and so require a longer and more thorough investigation. This distinction is important to take into account. Another aggravating factor is the fact that SpaceX aimed at human space-flight certification with the Falcon 9 vehicle, which it eventually reached. This might have biased the investigation length upwards, like it was observed in the Shuttle case.

Little information is known about the composition of these investigation committees, since seldom the full reports are published [52]. It was found that the Ariane 501 Inquiry Board was made up of 9 members, supported by a 5 member Technical Advisory Commitee [17]. The Ariane V142 flight Inquiry Board was made up of 7 members [21]. As for the SpaceX flights, the booster B1018 accident was investigated by about 12 members, whereas the B1028 accident team was closer to 20 members [64, 65].

Assuming a 15 member Board, with a yearly cost of 220 k€ per year (2021) [97], this translates into a cost of 253.151 k€ for a 28 day investigation, which corresponds to the ELV panorama, and a 1084.932 k€ for a 120 day investigation, analog to a reusable launch vehicle investigation duration.

$$C_{investigation_{ELV}} = 253.151 \text{ k€ (2021)} \tag{4.30}$$

$$C_{investigation_{RLV}} = 1084.932 \text{ k€ (2021)} \tag{4.31}$$

4.7.4. Implementation of changes

The change in production programs caused for instance by cost and performance improvements, are estimated as "delta development". These can be modeled as a set-back on the learning curve, or as additional development costs used for team training and production adaptation, in order to keep the rank on the learning curve [160].

Historical data on launchers has not shown a set-back on the learning curve [161]. According to the SPICE Reference Manual [161], the design and development (DD) factor can range between 0.1 and 0.2 for a minor modification, 0.5 for an average effort modification and 1 for a major modification with full qualification. In terms of testing and hardware models, a minor modification could be implemented on the PFM, an average modification could require a QM in addition to a flight model, and finally a major modification would require a full flight model and qualification model. These relate to the theoretical first unit according to a system test hardware (STH) factor, similarly to the model philosophy introduced in previous sections. In summary, this means that depending on the severity of the change needed to correct the failure, short of a new design or new development effort, the added cost could range between 1.1 and 3.3 of the T1 of the affected equipment. The determination of the appropriate factors are detailed in tables 4.8, 4.9 and 4.10.

Equipment Level DD	
DD effort	Factor
Minor Modification	0.1-0.2
Modification	0.5
Major Modification	1

Table 4.8: DD factors for delta development [161].

Scope of Design Effort	Models
Minor Modification	PFM
Modification	PFM or QM+PFM
Major Modification	QM+FM

Table 4.9: Determination of HW models [161].

STH Factors		
Model	Factor	Remark
FM	1.0	T1 reference
PFM testing	0.0-0.3	Avionics, dependent on amount of qualification
	0.0-0.5	Mechanical, dependent on amount of qualification
QM	1.1	Avionics product
	1.3	Mechanical product

Table 4.10: STH factors [161].

To more accurately model the cost of modifications along the life-time of a launch vehicle, there are two parameters that must be determined: First is the severity of the failure, and second is the source of the failure.

In terms of severity, it is usually assumed in most FMECA analysis that the relationship between the severity of a failure and its frequency is a negative exponential [90]. This implies that with a detailed FMECA analysis, it is expected that a designer can adjust the development and test hardware factors as a function the failure occurrence.

Regarding the source of failure, multiple studies have sought to determine the source of failures in launch vehicles [80, 129, 132]. These studies consistently show that failures attributed to the propulsion system account for the majority of accidents (>50%), followed by guidance and navigation systems (~20%). The expected provenance frequency of failures can be used in an analysis in order to achieve an estimate rooted in the source of the failure.

Having chosen the appropriate factors, the cost of modifications is obtained from the first unit with:

$$C_{modifications} = T1 \cdot (DD + STH + FM) \tag{4.32}$$

4.7.5. Cost of Downtime

The most extreme case of downtime after a failure is that of the Space Transportation System, which had an average number of days between consecutive flights of 82, and ran long investigations after its two accidents that resulted in down times of 975 and 922 days in 1986 and 2003 [61]. This long down time period and extensive failure investigation is intrinsically linked to the fact that the Shuttle was a manned system.

The Ariane family of launch vehicles has a long history; looking at the data from Ariane 4, the most flown launcher of the family, a downtime between 4 and 5 months can be consistently observed after each accident, with the launcher having on average one launch every 47 days [20]. It is important to safeguard however that intervals with no activity of similar length are not uncommon. With intervals of 4 months in at least 6 other occasions, these could have occurred due to a number of extrinsic factors, related to the development of Ariane 5 and alternating with other launchers from the previous and the next generation, at the start and end of the Ariane 4 era. The Ariane 5 launcher has only ever had two failures, the first on its very first launch, and an early failure on its 14th flight, where no significant downtime is recorded. There were 3 other instances of partial failures.

Looking at the SpaceX's Falcon 9 commercial history, the launcher boasts an average interval between flights of just 22 days [45]. The only two launch failures occurred in 2015 and 2016 and anticipated the two longest periods with no flights, at 177 and 133 days. Failures in the recovery of the first stage, which were more frequent, didn't observe a noticeable downtime.

Parkinson [150] estimates a cost relation to keep facilities running during a period of downtime. Based on TRANSCOST and ESA sources, it is established that it costs 72€ (2021) per kilogram per launch for remotely assembled vehicles in Western countries and 214€/kg/launch (2021) for vehicles assembled on the launch pad to maintain facilities and launch capability, considering the take-off mass of the vehicle. This is a figure that only takes into account the assembly characteristics of the vehicle, but not other properties such as type of propulsion (cryogenic propellants would be more expensive to store than storable propellants). This is therefore the best approximation possible with the data available.

Although the loss of business due to distrust in the launcher from the customer point of view is hard to quantify, the loss of revenue from a state of "Business as Usual" follows from the profit projection during downtime. Assuming a profit of 8% over the cost per flight, this cost figure will depend on the flight rate, on the time wherein launch capability is offline, and on the cost per flight.

$$C_{downtime} = C_{facilities} + C_{profit\ loss} \tag{4.33}$$

$$C_{facilities} = f_{dt} M \cdot \frac{L}{365} \cdot \#\ \textit{days downtime} \tag{4.34}$$

$$C_{profit\ loss} = 0.08 Cpf \cdot \#\ \textit{missed flights} \tag{4.35}$$

4.7.6. Summary

The model for the failure cost is summarized in table 4.11

Table 4.11: Failure cost model summary.

Failure Cost Model		
Element	Cost	
	ELV	RLV
Vehicle Replacement	$-$	MAN_{rs_i}
Flight Replacement	CpF_{avg}	
Insurance Penalty	$0.4CpF$	
Failure Investigation	253.151 k€	1084.932 k€
Modifications	$T1 \cdot (DD + STH + FM)$	
Downtime	$(72M + 0.08CpF) \cdot \#$ Flights missed	

The total failure cost can be obtained by summing all components:

$$C_f(T) = C_{replacement} + C_{insurance} + C_{investigation} + C_{modifications} + C_{downtime} \tag{4.36}$$

4.7.7. Validation

The validation of the failure cost model has been performed with a case study using data pertaining to the methane experimental (expendable) launch vehicle featured SOLSTICE [99].

Four scenarios are considered: a failure in the avionics subsystem, requiring minor improvements, a failure in the same subsystem, requiring a major modification, and failures in the engines, one requiring minor modifications and the other requiring major modifications.

The variable inputs used in the model are as follows:

- Cost per flight without development amortization: 39 M€;
- Amortized development factor: 4.37 M€;
- Theoretical first unit cost: Avionics - 2.41 M€; Engines - 2.44 M€;
- Lift-Off mass: 86.69 t;
- Flight rate: 8 launches/year;
- Insurance rate: 4%;
- Insurance bounce back period: 20 days;
- Investigation duration: 28 days;
- Two flights missed;
- Investigation board size: 15;
- Remotely assembled vehicle;
- Profit margin of 8%.

The results of this application can be seen on table 4.12.

Results from the expendable case validation show that the failure cost relative to the cost per flight, considering these subsystems and a best/worst case scenario analysis, falls within the recommendations prescribed in literature, where the failure cost was estimated to be around 2-5 times the cost of the vehicle, according to Koelle and Parkinson [129, 150].

Case	Avionics (Minor)	Avionics (Major)	Engine (Minor)	Engine (Major)
Re-flight	39.1			
Insurance	15.64			
Investigation	0.25			
Modifications	2.65	7.47	2.68	8.05
Facilities	62.42			
Profit Loss	6.26			
Total	126.31	131.13	126.35	131.71
Total relative to CpF	2.91	3.02	3.91	3.03

Table 4.12: Failure cost results for expendable launch vehicle (M€).

Considering that the first stage of the same vehicle could be reused, amortizing the cost of manufacturing the first stage over 8 flights, and that the cost of operations is twice that of the expendable case, the previous exercise is repeated. The inputs are as follows:

- Cost per flight without development amortization: 39 M€;
- Amortized development factor: 4.37 M€;
- Theoretical first unit cost: Avionics - 2.41 M€; Engines - 2.44 M€;
- First stage average production cost: 23.68 M€;
- Lift-Off mass: 86.69 t;
- Flight rate: 8 launches/year;
- Insurance rate: 4%;
- Insurance bounce back period: 20 days;

- Investigation duration: 120 days;
- Two flights missed;
- Four potential re-flights of reusable hardware lost.
- Investigation board size: 15 elements;
- Remotely assembled vehicle;
- Profit margin of 8%.

The result can be found in table 4.13.

Case	Avionics (Minor)	Avionics (Major)	Engine (Minor)	Engine (Major)
Re-flight	62.68			
Insurance	24.172			
Investigation	1.08			
Modifications	2.65	7.47	2.68	8.05
Facilities	62.42			
Profit Loss	7.77			
Total	160.77	165.59	160.81	166.18
Total relative to CpF	3.71	3.82	3.71	3.83

Table 4.13: Failure cost results for methane launch vehicle (reusable) (M€).

These results are also consistent with the recommendations found in literature. It is deemed therefore that a more detailed and accurate model has been developed that provides an estimate for the failure cost which falls within the results expected.

4.8. Expected Cost of Failure

The life-cycle reliability modeling can be used in conjunction with the estimated failure cost obtained with equation 4.36 in order to refine the life-cycle cost model.

It is common in civil engineering [114, 130, 173, 177] to estimate an expected cost of failure in the framework of calculating total life-cycle costs. This provides a measure of the cost of failure accumulated throughout the life-cycle, given the estimated failure cost and the probability of failure. This technique has also been applied in the aeronautic industry [189].

In this case, at any given flight, the expected cost of failure $C_F(T)$ is obtained with the expression:

$$C_F(T) = C_f(T) \cdot (1 - R(T)) \tag{4.37}$$

Where $C_f(T)$ denotes the estimate of the failure cost and $R(T)$ is the reliability of flight T.

This ”instantaneous” expected cost of failure translates into a measure of how much capital is being risked at flight T. However, the most useful measure concerns the expected cost of failure over a specific interval of time. This allows to extrapolate economic considerations, for instance, the expected cost of failure in the first years of operation or over the life-cycle of the launch vehicle. The analysis of how this compares to the cost per flight, as well as the revenue coming from these flights is an important measure of the value of the launch vehicle.

$$C_{F_{LCC}} = \sum_{i=1}^{n} [C_f(i) \cdot (1 - R(i))] \tag{4.38}$$

Where i denotes the flight number, and n the total number of flights in the life-cycle, obtained from the launch rate LpY multiplied by the number of years in the life-cycle of the launch vehicle.

Furthermore, for reusable launch vehicles, as was described in section 4.7.5, no downtime was noticed for failures related to landing. The first stage would necessarily need to be replaced, and it would be wise to assume that engineering modifications would be done in order to address failure modes that might have not been identified previously. On top of that, if the stage was insured, insurance premiums would be more expensive from the climbing insurance rates. However, it is possible that, since a failure in landing does not cause a loss of mission, no cost penalties are derived from the investigation and downtime domains. For that, it would be necessary to have two measures of probability: Firstly the probability of losing the mission, which necessarily results in loss of profits and accumulation of facilities cost from downtime, and would trigger a lengthy formal investigation. This corresponds to the unreliability of the launcher to accomplish the mission without catastrophically failing, as detailed in the previous chapter. The second measure of probability is the probability of having to replace the first stage, execute modifications and pay more insurance, which happens whenever the first stage is lost, due to a failure in accomplishing the mission or a failure in recovery. Considering the following events:

- A - Failure to accomplish the mission;
- B - Failure to accomplish the landing;
- C - Losing the first stage.

It is here assumed that a mission failure results in a total loss of the launcher and payload. The potential mission abort scenarios are not contemplated here for two reasons: Firstly, because in the Falcon 9 flight history there has never been a flight abort, and secondly, because the scenarios tested for human spaceflight lead to intentional destruction of the launcher, effectively reflecting this assumption. On the other hand, ”partial failures” where the payload is inserted into a wrong orbit are not considered, as the analysis focuses on launcher survival.

Losing the first stage happens therefore when either the mission is failed or the landing is failed. Using basic probability theory equations, it is possible to find P(C), the second measure of probability needed:

$$P(C) = P(A \cup B) = P(A) + P(B) - P(A \cap B) \tag{4.39}$$

In this context, the probability of failing the mission and the landing would refer to the partial failure situation. Since the analysis is focused on launcher survival, this probability is zero, as a failure in the mission leads to the impossibility of landing, and the resulting equation is:

$$P(C) = P(A \cup B) = P(A) + P(B) \tag{4.40}$$

With P(A) being equal to the unreliability of accomplishing the mission, or $1 - R$, and P(B) being the unreliability of the landing, or $1 - R_L$. This expression can therefore be further simplified:

$$P(C) = (1 - R) + (1 - R_L) = 2 - R - R_L \tag{4.41}$$

The expected cost of failure then becomes:

$$C_{F_r} = (1-R)(C_{Downtime} + C_{Investigation} + C_{Replace_{S2}}) + (2-R-R_L)(C_{Insurance} + C_{Replace_{S1}} + C_{Mod}) \tag{4.42}$$

In case one only has access to operational data, and considers a population consisting of several landing attempts, marking the failures, then the reliability obtained R_L' corresponds to the probability of accomplishing the landing, knowing that the mission was accomplished, or P($\overline{B}|\overline{A}$). Using the conditional probability properties this can be further decomposed:

$$P(\overline{B}|\overline{A}) = \frac{P(\overline{B} \cap \overline{A})}{P(\overline{A})} \tag{4.43}$$

As there is no intersection between events A and B, as shown in equation 4.40, then the numerator in 4.43 simplifies to $1 - P(A) - P(B)$. The true probability of failing the landing can then be found by solving to P(B):

$$P(B) = 1 - P(A) - P(\overline{A})P(\overline{B}|\overline{A}) \tag{4.44}$$

As we know, P(A) is the probability of failure of the mission, or $1 - R$, and the conditional probability term is the landing success probability taken from operational history, or R_L'. The expression can therefore take its final form:

$$P(B) = 1 - (1 - R) - R \cdot R'_L = R(1 - R'_L) \tag{4.45}$$

And P(B) was defined as the probability of failing the landing, or (1-R_L). We can therefore obtain the actual reliability of the landing, R_L:

$$R_L = 1 - R(1 - R'_L) \tag{4.46}$$

The probability of losing the first stage, $P(C)$, as a function of the mission reliability R and the probability of landing the stage given that the mission was successful, R'_L can therefore be computed:

$$P(C) = P(A) + P(B) = (1 - R) + R(1 - R'_L) = 1 - R \cdot R'_L \tag{4.47}$$

The cost of failure for reusable launch vehicles can be rewritten as a function of these two parameters:

$$C_{F_r} = (1-R)(C_{Downtime} + C_{Investigation} + C_{Replace_{S2}}) + (1-R \cdot R'_L)(C_{Insurance} + C_{Replace_{S1}} + C_{Mod}) \tag{4.48}$$

4.9. Expected Value

Given a discrete random variable X with frequency function $p(x)$, the expected value of X is given by the sum of the possible outcomes x_i weighted by their probabilities, $p(x_i)$ [119, 163]:

$$EV(X) = \sum_i x_i p(x_i) \tag{4.49}$$

From a cost perspective, measuring the expected value of a decision allows the designer to assess its economic feasibility [182], or determine preference [123, 124].

Value models have been used in the aerospace sector in order to measure economic parameters such as utility and risk [85]. Keller [123, 124] developed a space launch system value model in order to trade-off different architectures of launchers for a lunar Helium-3 mining mission. This model incorporated into the expected value computation cost and reliability, as well as a number of performance parameters. From a more cost-focused perspective, two works stand out: Firstly, Wagner analyzed the expected value of a vehicle payload escape system [182] in order to determine the economic feasibility of this system. Furthermore, Gülgönül and Sözbir apply a value model to a satellite project, weighting several costs, including the costs of the satellite and the launch, the insurance premium, and the revenue over lifetime, against the satellite's probability of survival [198].

In the context of estimating the value of a launch vehicle, the outcomes analyzed concern the profits or losses that result from each possible scenario in a launch.

In the case of an expendable launch vehicle, for a certain flight, the expected value is simply sum of the profit multiplied by the vehicle's reliability, and the outcome of a failure, multiplied by the vehicle's unreliability:

$$EV_e = R(PpF - CpF) + (1 - R)(PpF - CpF - C_f) \tag{4.50}$$

In the reusable case, the computation has to take into account the reliability of recovery or landing. In that sense, in case the mission fails, which has probability of 1-R, the cost of failure includes all the factors discussed in the failure cost section (minus insurance). However, if only the recovery fails, as was discussed in that section, the expected losses include only the replacement of the flight and some modifications.

$$\begin{aligned} EV = (1 - R)((PpF - CpF) - C_{downtime} - C_{investigation} - C_{Replace_{S1}} - C_{modifications}) \\ +(1 - R_L)((PpF - CpF) - C_{Replace_{S1}} - C_{modifications}) \\ +(R + R_L - 1)(PpF - CpF) \end{aligned} \tag{4.51}$$

Or, replacing R_L by the expression derived in equation 4.46 in terms of the experimental reliability of landing R'_L, the expression becomes:

$$EV = (1-R)((PpF - CpF) - C_{downtime} - C_{investigation} - C_{Replace_{S1}} - C_{modifications}) \\ +(R - RR'_L)((PpF - CpF) - C_{Replace_{S1}} - C_{modifications}) \\ +(RR'_L)(PpF - CpF) \tag{4.52}$$

A negative expected value would result in a unfavorable situation, where the launch service provider is expecting to lose money with each flight where EV<0, whereas a positive expected value results in an overall favorable situation despite the result being a success or a failure.

From this basic model, it is possible to model three more cases. These are: a situation where there is a Re-Fight Guarantee (RFG), where the first stage is insured, and a combination of both.

In case an RFG is agreed, then it is necessary to account in the simulation for an extra second stage, and the operations cost to re-launch the payload. Furthermore, besides the profit, there is also a premium at rate ins.

$$EV_{RFG} = (1-R)((PpF - CpF + ins \cdot CpF) - C_{downtime} - C_{investigation} - C_{Replace_{S1}} - C_{Replace_{S2}} - C_{OPS} \\ -C_{modifications}) \\ +(R - RR'_L)((PpF - CpF + ins \cdot CpF) - C_{Replace_{S1}} - C_{modifications}) \\ +(RR'_L)(PpF - CpF + ins \cdot CpF) \tag{4.53}$$

If the first stage is insured, the cost of replacing the first stage is covered, but it is necessary to pay the premium for that insurance, at rate ins:

$$EV_{ins} = (1-R)((PpF - CpF - ins \cdot MAN_{rs}) - C_{downtime} - C_{investigation} - C_{modifications} - C_{insurance}) \\ +(R - RR'_L)((PpF - CpF - ins \cdot MAN_{rs}) - C_{modifications} - C_{insurance}) \\ +(RR'_L)(PpF - CpF - ins \cdot MAN_{rs}) \tag{4.54}$$

The combination of these two situations is modeled as follows:

$$EV_{RFG,ins} = (1-R)((PpF - CpF + ins(CpF - MAN_{rs})) - C_{Replace_{S2}} - C_{OPS} - C_{downtime} - C_{investigation} \\ -C_{modifications} - C_{insurance}) \\ +(R - RR'_L)((PpF - CpF + ins(CpF - MAN_{rs})) - C_{modifications} - C_{insurance}) \\ +(RR'_L)((PpF - CpF + ins(CpF - MAN_{rs}))) \tag{4.55}$$

An application of this model will be worked out in chapter 5.

On top of this, it is important to take into account that utility isn't quantified by the result of the expected value [76]. A risk-adverse company would base its decisions on the outcome of the expected value computation, whereas a risk proverse one could find opportunity in parts of the life-cycle where the expected value is lower, given the relatively low risk of failing the mission or landing versus the high payoff. Arrow and Lind [79] find that a government should take a risk neutral stance, and so a governmental initiative such as the SLS would opt for decisions that lead to the greatest expected value [124]. It follows then that a private company such as SpaceX could in theory adopt a more risk adverse or risk proverse stance. This would imply the application of utility theory to the results obtained from expected value. It is considered that this falls out of the scope of this research, however it constitutes an important caveat when analyzing the outcome of the expected value computation, and taking into account the overall probability of failure versus its outcomes.

4.10. Summary and Conclusion

In this section, the overall cost methodology was outlined, and the cost-reliability link was established, both at the reliability increase through redundancy and testing in the design phase, and at the operational phase with the cost of failure.

The SOLSTICE T1-equivalents method was described and modified to suit the case at hand. The operations methodology described in TRANSCOST was appropriated and applied to the reusable launch vehicles as they operate in the current market landscape. To these models, a set of CERs from NASA CASTS was repurposed and added in order to estimate the cost of reusable launch vehicle hardware. These were calibrated with adjustment factors derived from ESA project mass and cost data. The refurbishment costs also follow from the adaptation of a NASA CER, coming from the OCM. This was modified with an aging factor, as previously applied by Rozemeijer. The retrieval costs were included, by deriving equations suitable to estimate the costs of modern recovery strategies, developed in a bottom-up fashion by DLR.

In addition to this, a failure cost model was developed, based on recommendations present in TRANSCOST, historical data and ESA processes. This model was validated against the initial expectations, and for the first time an estimate for the failure cost is achieved.

Finally, the reliability model is integrated with the cost estimates, and expressions for the expected cost of failure and expected value are derived, in order to be applied to the reusable launch vehicle scenario.

The model outlined allows designers to consider the failure as a cost parameter, expressed as the expected cost of failure. This allows for a better refining of results when comparing different launch vehicle architectures and solutions, as the reliability of the design has a direct impact in terms of cost. The expected value model further facilitates this trade-off, by allowing to directly compare the value of different solutions according to a flight’s cost, reliability, and outcome scenarios. In future optimization studies, this model can be integrated in order to refine results by adding reliability as a variable that expresses a direct effect on the cost per flight.

5
Falcon 9 Case Study

The primary case study for this research is naturally SpaceX's Falcon 9 launcher, as it is the only reusable launch vehicle to have ever been successfully used to deliver payloads to orbit. In this chapter, the cost model outlined in the previous chapter is applied to the current version of Falcon 9, according to the most up to date information publicly available. The cost estimates obtained are validated along the process, against multiple sources, from parallel estimates to information provided by SpaceX officials in interviews.

In parallel, a reliability estimate is obtained, based on operational data of Falcon 9 and historical data from comparable launchers.

By obtaining solid and credible cost estimates, the scenario obtained from the final value model, which is dependent on more uncertain reliability estimates, will be lent a higher degree of credibility. This scenario will be challenged in a sensitivity analysis study, in order to judge the impact of parameter uncertainty on the final result.

5.1. System Configuration

The Falcon 9 launcher is a two-stage reusable launch vehicle, with a payload capacity of 22,800 kg to LEO and 8,300 kg to GTO [66]. The first stage is propelled by 9 sea-level Merlin engines, while the second stage features only one vacuum Merlin engine.

In terms of hardware related to reusability, the Falcon 9 launcher features:

- 4 landing legs;
- 4 hypersonic grid fins;

It should be mentioned that although previous iterations of the launcher featured an ablative heat shield at the bottom of the first stage in the engine cluster (from Falcon 1 up until Block 5) [33], there is no report of a thermal protection system being employed. Some reports mention the use of Inconel [33, 34] as a replacement, which lead to savings in the refurbishment cost. Furthermore, in declarations by Elon Musk as well as in literature [34, 195] it is determined that thermal protection is not necessary.

5.2. System Mass

The main challenge related to providing a cost estimate using T1 equivalents for this case is the lack of concrete mass data about the launcher systems. This is due to SpaceX being a private company, which results in a larger degree of opacity related to the specifics of the launcher, when compared to NASA and ESA programs. A number of assumptions related to the systems characteristics will be needed in order to develop an estimate with this method.

Firstly, the publicly available information related to the Falcon 9 vehicle is gathered [4]:

- Mass: 549,054 kg;
- First Stage Dry Mass: 25,600 kg;

- First Stage Propellant Mass: 395,700 kg;
- Second Stage Dry Mass: 3,900 kg;
- Second Stage Propellant Mass: 92,670 kg;

This information can be used in conjunction with the model by Rozemeijer [165] to estimate the mass of some of the recovery and landing hardware.

The mass of the grid fin can be obtained from its frontal area following the relationship [165]:

$$m_{fin} = 26.028 S_{fin} \tag{5.1}$$

With this relation, the mass estimate for a single grid fin is 62kg. This estimate presents a 6.2% error compared to the MER result of 58kg obtained from the Space Shuttle experience when using the same area value [164].

Assuming a thickness of 0.01 m [155], and knowing the density of titanium $\rho_{Ti} = 4420\ kg/m^3$, this would correspond to a Titanium plate with 41% of the material removed.

The cost of the grid fin system is estimated using the CASTS T1 CER for Mechanisms, which is suited to cost miscellaneous mechanical systems [145]. An adjustment factor derived from ESA data related to analogous system is used to calibrate the result.

The mass of the landing legs is assumed to be 10% of the mass of the recovered system [165, 174].

$$m_{legs} = 0.1 m_{S1} \tag{5.2}$$

This of course would vary with the amount of unused propellant. In this case, in order to simplify calculations, it shall be assumed that there is no reserve propellant left in the tanks after landing. This isn't considered a driving cost factor in this study, as it only influences the cost of the recovery system, which is just a fraction of the total cost of the launch vehicle. However, when inserted in an optimization environment, it is advised to consider this variation of the landing gear mass with the total landing mass including consumables.

Assuming that upon landing all propellant has been expended, this leads to a mass estimate of 640 kg per leg. The cost estimate for this subsystem is obtained using the CASTS T1 CER for recovery systems [145], using an adjustment factor obtained from ESA data of comparable systems in order to calibrate the result.

The remaining total mass breakdown for the Falcon 9 reusable launcher was obtained using four different resources and is presented in table 5.1. Their cost estimates will be obtained using the CERs from SOLSTICE.

- Publicly available information and data [6, 35];
- Internal ESA cost Falcon 9 cost estimates using SPICE and SOLSTICE, using a mass breakdown which is itself based on publicly available data (less at the time) and the Ariane 4 mass breakdown, which was considered an analogous system to Falcon 1 given its configuration [98, 99, 162];
- Verification of order of magnitude with mass estimating relationships, whenever needed/appropriate [75, 164];
- Expert judgment [97].

Element	Unit	Equipment	Part	Mass (kg)	Number
Stage 2	Payload	Payload Adapter		100	1
		Payload Fairing		1,800	1
		Sub-Total		1,900	
	Stage	Pressurizant Tank		81	1
		Propellant Tank		646	1
		Stage Structure	Thrust Cone	190	1
			Skirt	574	1
		Engine(s)		470	1
		Thrust Vector Control		94	1
		Pipes Valves	Pipes	50	1
			Valves	30	1
		Stage Harness		25	1
		Sub-Total		2,160	
	Avionics	Avionics	Comms	2	2
			Power	12	2
			Data Handling	19	2
			GNC	6	2
			Avionics Harness	11	2
		Sub-Total		99	
	Attitude Control	Attitude Control Module		27	1
	Stage Total			4,186	
Stage 1	Interstage	Interstage Structure		814	1
	Stage	Pressurizant Tank		243	1
		Propellant Tank		4,742	1
		Stage Structure	Thrust Cone	2,060	1
		Engine(s)		470	9
		Thrust Vector Control		196	1
		Pressurization System		134	1
		Pipes Valves	Pipes	487	1
			Valves	70	1
		Stage Harness		40	1
		Sub-Total		12,202	
	Avionics	Avionics	Comms	2	2
			Power	18	2
			Data Handling	29	2
			GNC	9	2
			Avionics Harness	16	2
		Sub-Total		149	
	Attitude Control	Attitude Control Module		41	1
	Recovery	VTVL	Landing Gear	640	4
			Grid Fins	62	4
		Sub-Total		2,810	
	Stage Total			16,016	
Launch Vehicle Total				20,202	

Table 5.1: Falcon 9 Mass Breakdown.

5.3. Cost estimation

In this section, the cost model outlined in the previous chapter is applied to the Falcon 9, according to the most recent publicly available data. The estimates are validated against information available.

5.3.1. Theoretical First Unit Cost

The theoretical first unit costs (T1) of the Falcon 9 launcher was obtained following the methodology outlined in the previous chapter. The T1 estimates for the recovery system (landing legs and grid fins) were obtained using the CASTS CERs, calibrated with the adjustment factors presented in chapter 4, section 4.2.1. The T1 estimates of the remaining systems were obtained with the SOLSTICE CERs [99]. The results are summarized in table 5.2.

Element	Unit	Equipment	Part	T1 (M€)
Stage 2	Payload	Payload Adapter		0.57
		Payload Fairing		6.21
		Sub-Total		6.78
	Stage	Pressurizant Tank		0.50
		Propellant Tank		2.19
		Stage Structure	Thrust Cone	0.36
			Skirt	1.00
		Engine(s)		4.37
		Thrust Vector Control		0.59
		Pipes Valves	Pipes	0.14
			Valves	0.10
		Stage Harness		0.13
		Sub-Total		9.38
	Avionics	Avionics	Comms	0.08
			Power	0.33
			Data Handling	1.61
			GNC	0.34
			Avionics Harness	0.09
		Sub-Total		2.44
	Attitude Control	Attitude Control Module		1.60
	Stage Sub-Total			38.81
Stage 1	Interstage	Interstage Structure		0.70
	Stage	Pressurizant Tank		1.09
		Propellant Tank		9.05
		Stage Structure	Thrust Cone	3.21
		Engine(s)		4.37
		Thrust Vector Control		0.92
		Pressurization System		2.36
		Pipes Valves	Pipes	0.69
			Valves	0.18
		Stage Harness		0.15
		Sub-Total		22.02
	Avionics	Avionics	Comms	0.11
			Power	0.45
			Data Handling	2.22
			GNC	0.48
			Avionics Harness	0.10
		Sub-Total		3.36
	Attitude Control	Attitude Control Module		2.46
	Recovery	VTVL	Landing Gear	1.63
			Grid Fins	0.02
			Sub-Total	1.66
	Stage Sub-Total			57.24
Launch Vehicle Sub-Total				96.05
Integration and Testing				3.07
Launch Vehicle Total				99.12

Table 5.2: Falcon 9 T1 Results.

Recovery System Verification

Given the opaque nature of the CERs used in the calculation of the T1 estimate of the recovery system, the result obtained is verified with the model used in Rozemeijer and Contant's estimates [86, 165]. This methodology involves including the landing legs and grid fins in the estimate of the overall structural mass of the stage, and obtaining a cost estimate using the appropriate CER taken from TRANSCOST. These can be consulted in [165].

The factors used in the calculation can be found summarized in table 5.3. The rationale behind the selection of these factors can be found in [165].

Symbol	Parameter	Value
M_{rec}	Recovery Mass	16,016 kg
M_{rs}	Mass of Recovery System	13,208 kg
WYr	Work Year Costs (2021)	0.314 M€
f_8	Country Productivity Factor	0.86
f_{10}	Techn. Program Cost Factor	0.8
f_{11}	Commercialization Factor	0.5

Table 5.3: Input parameters for alternative model to estimate the cost of recovery system [86, 165].

The resulting T1 estimate of the recovery system obtained by implementing this alternative model is 1.617 M€. The estimate obtained with the present model, for both the landing legs and grid fins, was 1.66 M€. This corresponds to an error of 2.47% of the present model, when compared to the alternative model.

5.3.2. Development Cost

The development cost was estimated based on the T1, as outlined in the methodology present in 4.3. A summary of the parameters used can be found in table 5.4.

Parameter	Value
f_8	1
M/PA%	5.3%
DD	3-6
STH	3.1
c_p	0.97
I&T	0.032

Table 5.4: Development phase parameters.

A country productivity factor f_8 of 1 is chosen, as this is standard in a preliminary estimate, where the country of production is not specified. A management and product assurance percentage (M/PA%) of 5.3% is selected as recommended in SOLSTICE [99], for companies with few hardware subcontractors such as SpaceX. The design and development (DD) T1 equivalent selected for most components is 3, as suggested in SOLSTICE [99]. This is because to minimize costs, it is logical to assume that flight proven hardware is used, therefore incurring in no additional TRL increase at the development level. The DD factor of 3 indicates the TRL increase that is needed whenever a new component is integrated in a new system. This includes the engine, which is one of the main cost drivers at development level, since the Merlin engine was already flown on the Falcon 1 launch vehicle. The two exceptions are the first stage GNC and the recovery hardware systems. This is because it is assumed that extra development was needed in these elements, in order to achieve the landing capability. In the case of the GNC system, it is assumed that the technology was at a TRL of 3, which denotes that the element concept is elaborated and expected performance is demonstrated through analytical models supported by experimental data and characteristics. The DD factor for the GNC system is therefore 6. For the landing legs and grid fins, it is assumed that the original TRL is 5, denoting that the critical functions of the element are identified and the associated relevant environment is defined. This results in a DD factor of 4 [107].The reasoning for the system test hardware equivalent (STH) follows from the proto-flight model assumption [99], by considering a demonstration model equivalent to 30% of the T1, an engineering model equivalent to 130% of the T1 and combining the qualification model and flight model into a protoflight model, which including refurbishment costs is equivalent to 150% of the T1. The profit retention factor c_p is taken at 0.97, as recommended in SOLSTICE. Finally, the integration and testing effort considered is 3.2%, which is derived in SOLSTICE.

The results are summarized in table 5.5.

Element	Development Cost (M€)	Percentage of Total (%)
Stage 1	164.8	44.1
Stage 2	111.1	29.8
Engine	68.1	18.2
Recovery Hardware	17.8	4.8
I&T	11.6	3.1
Vehicle Total	373.3	100

Table 5.5: Falcon 9 Development Costs

As can be seen in table 5.5, the total development cost obtained for the Falcon 9 launch vehicle is 373.3M€ (FY2021). The cost to develop the recovery system corresponds to about 5% of the total development cost. On the other hand the engine's development cost is close to a fifth of the total development cost.

The development cost corresponds to $353M in the year 2010. The publicly available references concerning the Falcon 9 development costs include:

- NASA cost estimates done with NAFCOM, which resulted in an estimate of $4B (2010), when using a traditional NASA approach, and $1.7B (2010) when using a SpaceX more commercial approach [141, 143]. This original estimate was developed based on a mass estimate and technical data from SpaceX, and was modelled as a cost plus fee model. This estimate has since been updated, based on updated technical and mass corrections, plus insight into the parametrization of the estimate with a physical visit of the estimators to the SpaceX facilities. The updated estimates were $1.4B (2010) using a cost plus fee acquisition model, and $440M (2010) using a firm fixed price acquisition.

- Elon Musk's statement that total development cost for the launcher was $300M (2010) [36]. If including the development of the Falcon 1 vehicle, estimated at $90M (2010), the development cost of Falcon 9 would be maximized at $390M (2010) using this data, which was verified by NASA [141].

- Insight taken from the Falcon 9 funding program: Under the Commercial Orbital Transportation Services (COTS) agreement, SpaceX received $278M in 2006 and $118M in 2011, which correspond to a total of $372M (2010). The Falcon 1 vehicle development was funded privately.

The estimate obtained with this methodology presents a 19.7% error when compared to the NAFCOM estimate. Compared to the Falcon 9 funding scheme, the error is 5.1%. Compared to the figures provided by Elon Musk regarding the Falcon 9 development, the error obtained is 17.7%, or 9.5% when including the Falcon 1 development figure. The estimate is therefore within the 30% error standard considered for this exercise.

5.3.3. Manufacturing Cost

Manufacturing costs are obtaining using a 90% Crawford learning curve applied to the T1 costs.

It is assumed in this exercise that every first stage (or booster) has a useful life of 10 flights, after which it is expended. This follows from the observation that the maximum number of re-flights a booster has had is 12, and that there has not been registered any significant offline refurbishment on a booster, and so in the operations section, only online refurbishment shall be considered.

The simulated costs are depicted in figure 5.1 and summarized in table 5.6. It can be seen that the manufacturing cost of the first flight of the life-cycle of each booster is disproportionately more expensive, when compared to the rest. This happens because the whole first stage, including the interstage, the recovery hardware, and 9 Merlin engines need to be produced. This is a lot more expensive when compared to the remaining flights, where just the second stage, which only contains one Merlin engine, is produced.

Booster	1		2		3	
State	New	Reused (avg)	New	Reused (avg)	New	Reused (avg)
MAN	87.9	12.7	71.8	10.9	66.5	10.1

Booster	4		5		6	
State	New	Reused (avg)	New	Reused (avg)	New	Reused (avg)
MAN	63.3	9.6	61.0	9.3	59.2	9.0

Booster	7		8		9	
State	New	Reused (avg)	New	Reused (avg)	New	Reused (avg)
MAN	57.7	8.8	56.5	8.6	55.5	8.4

Booster	10	
State	New	Reused (avg)
MAN	54.6	8.3

Table 5.6: Summary of the manufacturing costs, in M€.

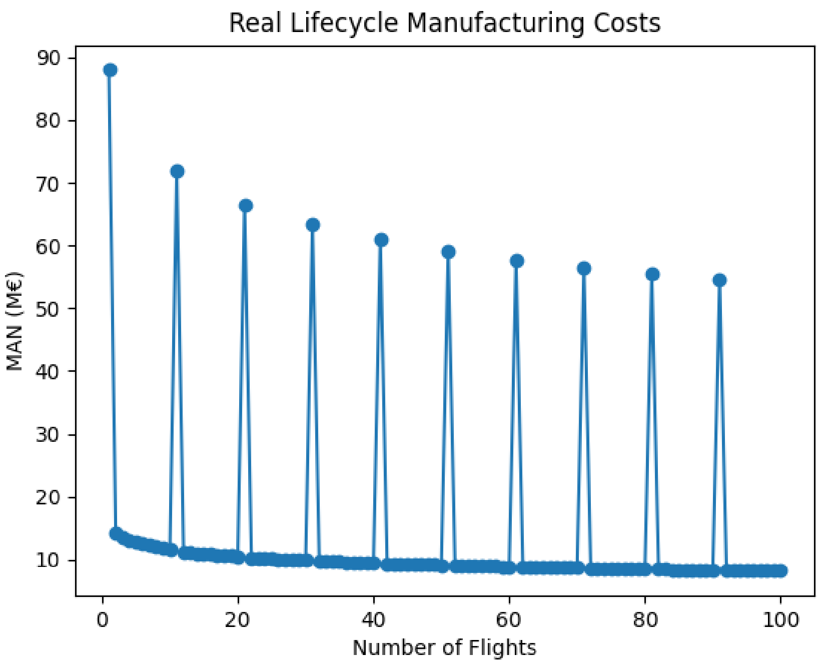

Figure 5.1: Falcon 9 real manufacturing costs.

By amortizing the cost manufacturing each new booster over the respective flights, it is possible to have a more accurate view of the manufacturing costs per flight. This is represented in figure 5.2. From the graph, it is possible to distinguish the effect of the learning occurring at engine level, as from a set of booster flights to the next a cost reduction can be noticed. This happens because producing 9 engines for each booster accelerates learning, and since the engine is a costly hardware element, a reduction in its contribution is noticeable.

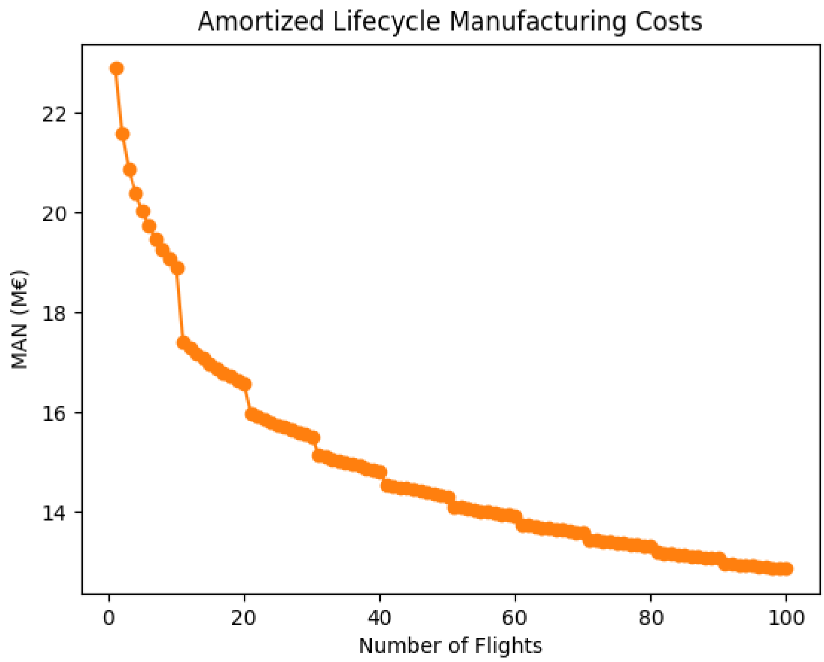

Figure 5.2: Falcon 9 manufacturing costs with amortized booster costs.

In figure 5.3, the different ways to visualize cost are compared. Both obviously have the same value when averaged out, which is 14.9M€ (2021).

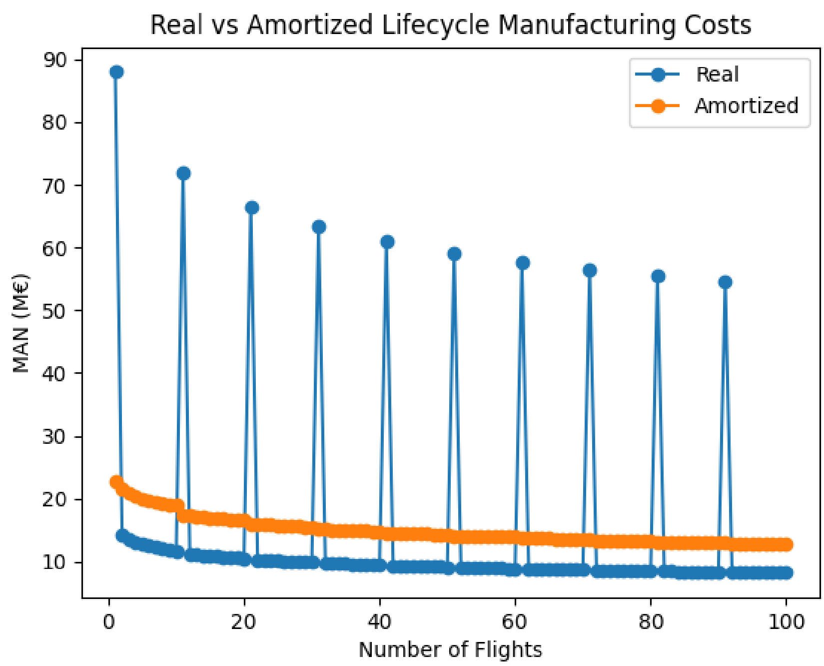

Figure 5.3: Falcon 9 manufacturing costs, comparing the real costs to the amortized version for better visualization.

The only reference available for the upper stage cost is that of an interview Elon Musk gave Aviation Week on May of 2020 [23]. There, he says that the manufacturing cost of a Falcon 9 upper stage is $10M. According to this methodology, the average cost of the upper stage over 100 flights was 9.6M€

(2021), which in 2020 dollars is $11M. This presents a 10% error related to the declarations (although it could be argued that the figure provided was not exact, and referred more to a rough estimate). It is however a good sign that this figure is within the standard 30% range.

5.3.4. Operations Cost

The operations cost follows from the methodology described in chapter 4. The launch rate considered is 20, as it is the average launch rate achieved by SpaceX since the debut of the Full Thrust version of Falcon 9 in December 2015. A return to launch site recovery method is selected. The most important point to safeguard here is that the refurbishment is considered to only happen at the engine level, following from evidence coming from SpaceX officials [8, 29].

The inputs are summarized in table 5.7.

Element	Symbol	Value	Unit
Learning Curve DOC	lc_{DOC}	0.85	-
Learning Curve REF	lc_{REF}	0.9	-
Aging Factor	f_a	1.15	-
Rate Curve	rc	0.5	-
Work-Year Cost	WYr	0.314339	M€
Payload Mass	M_{PL}	22800	kg
Propellant Mass	M_p	488370	kg
O/F ratio	r	2.56	-
Specific Cost RP-1	c_f	0.06528	$€/kg$
Specific Cost LOX	c_{ox}	0.15232	$€/kg$
Annual Flight Rate	LpY	20	-
Reusable HW lifespan	L	10	-
Number of Stages	N	2	-
Vehicle Type Factor	f_v	1	-
Assembly and Integration Mode Factor	f_c	0.7	-
Country Productivity Factor	f_8	1	-
Commercial operation multiple	f_{11}	0.55	-
Vehicle Complexity Factor	Q_N	1.4	-
Specific Transportation Cost	c_{TRA}	5.837	$€/kg$
Percent of work subcontracted out	σ	20	%
Payload site charge fee	c_{PL}	5.995	$€/kg$
Launch Site Fee	c_{site}	1.327	M€
Public Damage Insurance	I	0.1	M€

Table 5.7: Parameters used in calculation of Operations cost.

The most relevant results are summarized in table 5.8:

Element	Average Cost (M€)
DOC	8.35
IOC	3.90
REF	0.225

Table 5.8: Operation costs summary.

Once again comparing the refurbishment costs obtained to the reference provided by Elon Musk in his Aviation Week interview in May 2020 [23] of $250k (2010), using the 2010 dollar equivalent of $214k, an error of 14.4% is achieved. This is under the accepted 30% margin. Furthermore, according to guidelines provided by TRANSCOST, it is expected that for reusable launch vehicles, Direct Operations Costs amount to 45-60% of the total cost, Indirect Operations Cost amount to 10-15% and the Vehicle Recurring Cost (which includes manufacturing cost with amortization and refurbishment) would amount to 30-40% [129]. In this study, the DOC amounted to 30%, the IOC to 15%, and the VRC to 55%, as illustrated in figure 5.4.

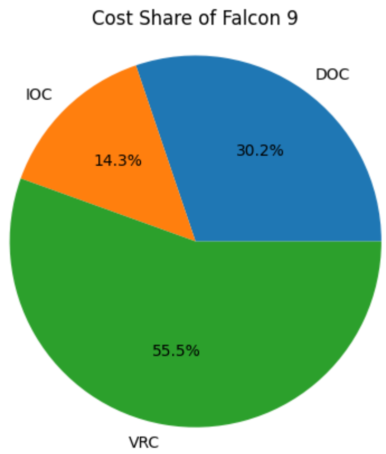

Figure 5.4: Falcon 9 cost share results according to the model.

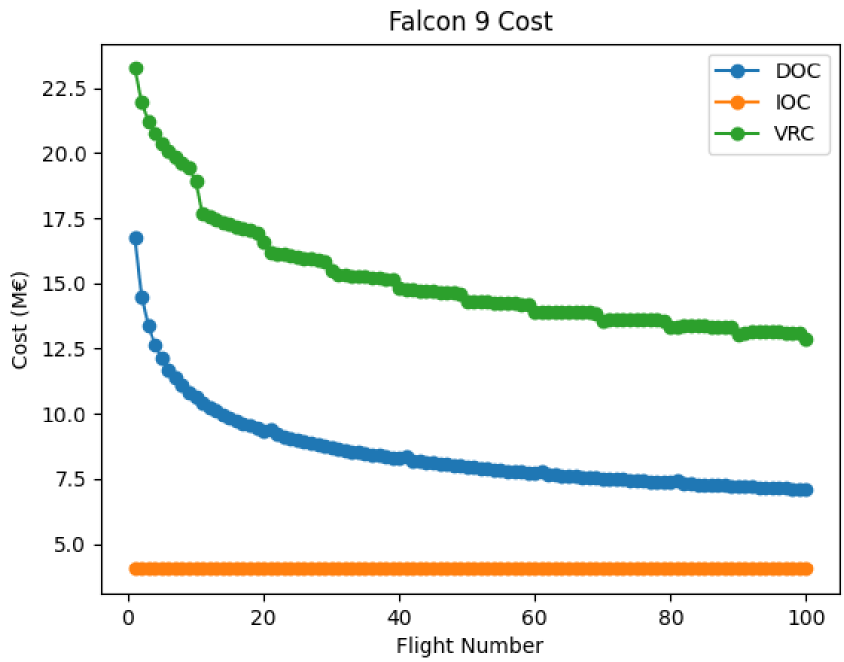

Figure 5.5: Falcon 9 operations cost with time.

As expected, the IOC fits the range given. However, the VRC is inflated by at least 15%, which should be attributed to the DOC. This is however expected: the TRANSCOST reference assumes that the reusable hardware is capable of flying for longer life spans than 10 flights. In this study, a purposeful decision to limit the reuse number to 10 was made, in order attempt to better represent the current reusable launch vehicle industry, which is slightly different than expected when TRANSCOST was written, heavily influenced by the Space Shuttle experience.

The variation of these costs throughout the life time of 100 flights considered is depicted in 5.5 and is summarized in table 5.9.

Booster	1	2	3	4	5
OPS (avg)	16.4	13.8	12.9	12.4	12.1

Booster	6	7	8	9	10
OPS (avg)	11.8	11.6	11.4	11.2	11.1

Table 5.9: Operations cost summary in M€.

5.3.5. Cost and Price per Flight

A reference from May of 2018 [28] contains a relative cost breakdown from Elon Musk pertaining to the Falcon 9 launch vehicle. He states that the booster is close to 60% of the cost, the upper stage 20%, the fairing 10%, and the final 10% are connected to the flight cost itself. Considering that by May 2018, the company had recorded 34 Full Thrust Falcon 9 launches, and had produced 28 Full Thrust boosters, we can extrapolate the learning curve rank of the different components: The first stage rank is 28, the upper stage is 34, and the engine rank is 286. The fairing rank is more difficult to track, but it is assumed that it stood at 34. In terms of operations, the rank is assumed at 65, to include all flights from Falcon 9 until this time, and the flight rate considered is 20 (the number of Falcon 9 launches in 2018). It should be noted that refurbishment and recovery costs were not included, as they are only applicable to reused versions of Falcon 9. Furthermore, the commercialization costs (IOC) were also left out, as they are not related to the cost itself, not being charged for instance for government contracts [99], or when SpaceX is using its own launchers to fly satellites of the Starlink constellation they own. The cost breakdown is as presented in figure 5.6:

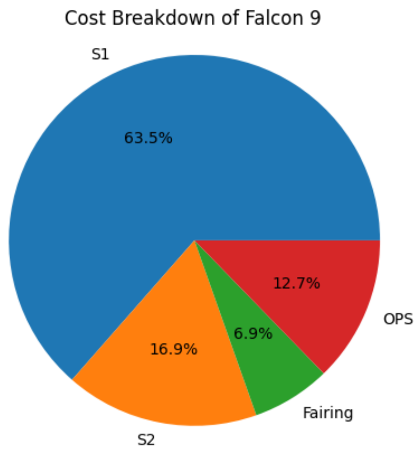

Figure 5.6: Falcon 9 estimated Cost Breakdown as of May 2018

It can be seen that all figures are within reasonable acceptable ranges (taking into account the nature of the reference), being accurate to the nearest tens range.

Adding the Manufacturing and Operating costs, and amortizing the Development costs over the expected number of flights over 10 years, as per an internal ESA reference, the Cost per Flight over the lifetime of the launcher is projected, and illustrated in figure 5.7. It should be noted that since a commercial offering is considered, these costs include the commercialization costs expressed as an indirect operating cost. Launches connected to military or NASA contracts, or Starlink launches will therefore feature a different cost. The relevant information is summarized in table 5.10, and a breakdown of costs is present in chart 5.8.

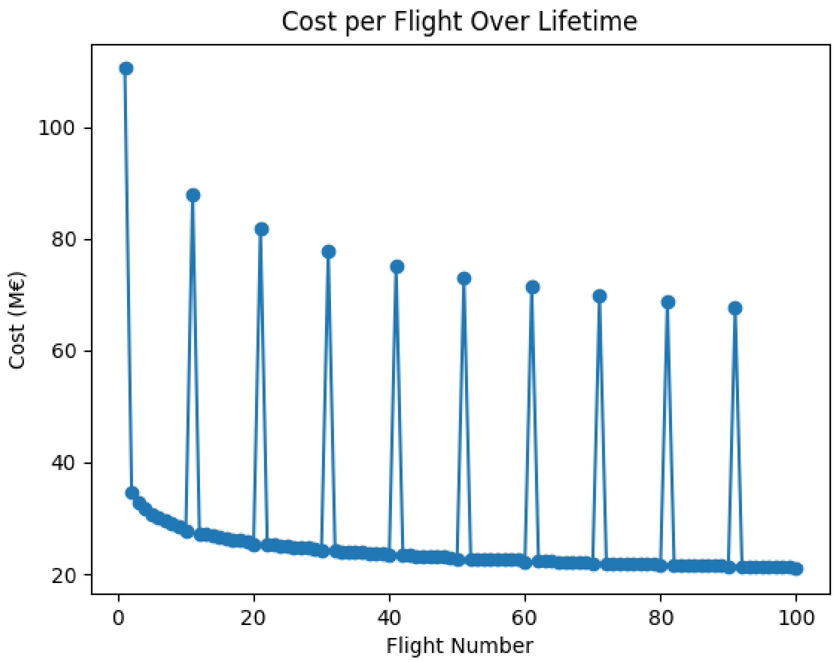

Figure 5.7: Falcon 9 Cost per Flight over Life Time

Flight Type	Average Cost (2021 M€)	Average Cost (2021 M$)
New Launcher	78.3	92.6
Launch w/ Reused Booster	23.8	28.1
Total	29.3	34.65

Table 5.10: Cost Per Flight result summary.

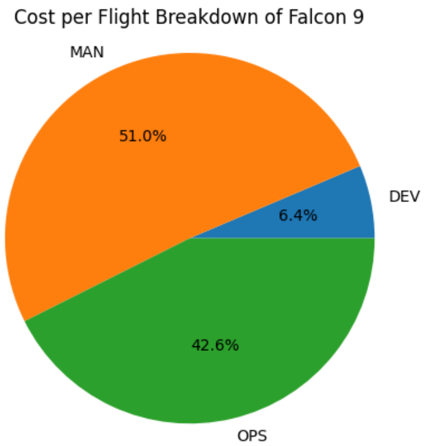

Figure 5.8: Falcon 9 estimated Cost per Flight Breakdown.

From these results, it is possible to see that the average cost of a launch using a reused booster is less than a third of the average cost of the flights that require building a new booster. This supports Elon Musk's assertion [30]) that there is a cost benefit starting from the third reuse on-wards.

The cost of a flight using a new booster is actually comparable to that of the Ariane 6: The A64 version has a payload capacity to LEO of 21,650 kg [32], close to Falcon 9's 22,800 kg [66], and has a projected price of 115M€ per flight. Removing ESA's standard profit margin of 8%, it is assumed that the CpF of Ariane 6 in this configuration would be around 106.5 M€. This would mean that the Falcon 9 vehicle in terms of cost outperforms Ariane 6, for the same mission type, in both its reused and brand new form.

Taking pricing into account, a Falcon 9 launch using a new booster before 2022 costs the customer $62M (at least since 2018 [25], until 2022 when it was raised to $67M, to cope with inflation [66]). Furthermore, it has been reported that opting for a flight using a reused booster would be priced at around $50M [31]. Comparing to the results obtained here, this would mean that the flights on new launchers were being sold at a loss (costing 49% more than they are sold for), whereas the much cheaper launch on reused boosters are being sold at a profit that is much higher than the 8% standard for ESA launches [99] (78% profit margin, if a price of $50M is considered). Considering a period being studied where 10 flights are flown with new boosters, and 90 flights with reused boosters, assuming a flat price per flight of 50 million dollars for flights with reused boosters and 62 million dollars for flights with a new booster, the average profit margin is 48%.

This would be a sound pricing strategy, as it entices commercial customers with an offer that actually results in a loss, in order to build up the reputation and flight history of the launcher, to later make higher profits on the flights using reused boosters.

There is however a caveat to this pricing analysis. The flights flown for governmental clients, such as NASA, DOD and NRO are priced much higher than the commercial offer, from 80 million dollars, to contracts worth more than 300 million dollars, for a Falcon Heavy launch which was priced at 97 million dollars for commercial customers [22, 26, 40, 42, 62]. On top of this, after June 2021, SpaceX was given approval to fly reused boosters for governmental contracts. All of this results in much higher profit margins in flights servicing these public clients. Furthermore, in one of these flights, SpaceX director of vehicle integration Christopher Couluris stated that the cost of the launch was $28M (2018). Comparing this to the average cost obtained of $34.65M, it represents an error of 23.2%.

5.3.6. Failure Cost

The inputs used in the computation of the failure cost in each flight are summarized in table 5.11. On every flight, it is assumed that a standard modification is needed, rather than a major or minor one, and that two thirds of the modifications are done on the engines and one third on the avionics systems. This comes from the observation that this is close to the failure cause distribution for launchers [129, 132].

Element	Symbol	Value	Unit
Insurance Rate	ins	4%	-
Insurance Recovery	n_i	20	days
Investigation Board Rate	W_{inv}	0.22	M€/person/year
Board Size	n_{inv}	15	person
Time of Investigation	t_{inv}	120	days
Design and Development Factor	DD_{mod}	0.5	-
Frequency of Engine Failures	-	66.6%	-
Frequency of Avionics Failures	-	33.3%	-
Protoflight Model Equivalent	PFM_{mod}	0.25	-
Facility Cost Factor	f_{dt}	72	€/kg/launch
Downtime	t_{dt}	155	days
Average Profit Margin	p	48%	-

Table 5.11: Failure cost inputs.

The results are displayed in the graph present in figure 5.9. It can be seen that the overall trend for the failure cost is decreasing, as some of its terms are proportional to the cost per flight or manufacturing cost which are themselves decreasing. It can also be observed that the last flight in each boosters life time is less expensive in case of failure. This is due to the fact that the on the last flight, the booster is expended or retire, and so a failure wouldn't require replacement.

The average failure cost throughout the lifetime of this vehicle is 494 M€, which is close to 17 times the cost per flight. The reason why this value is much higher than those predicted by Koelle and Parkinson [129, 150], who postulated that it should be between 2 and 5 times the cost per flight, is because the failure cost scales with the yearly launch rate. In TRANSCOST, usually a launch rate between five and eight flights per year is used. In this exercise, a launch rate of 20 flights per year is selected.

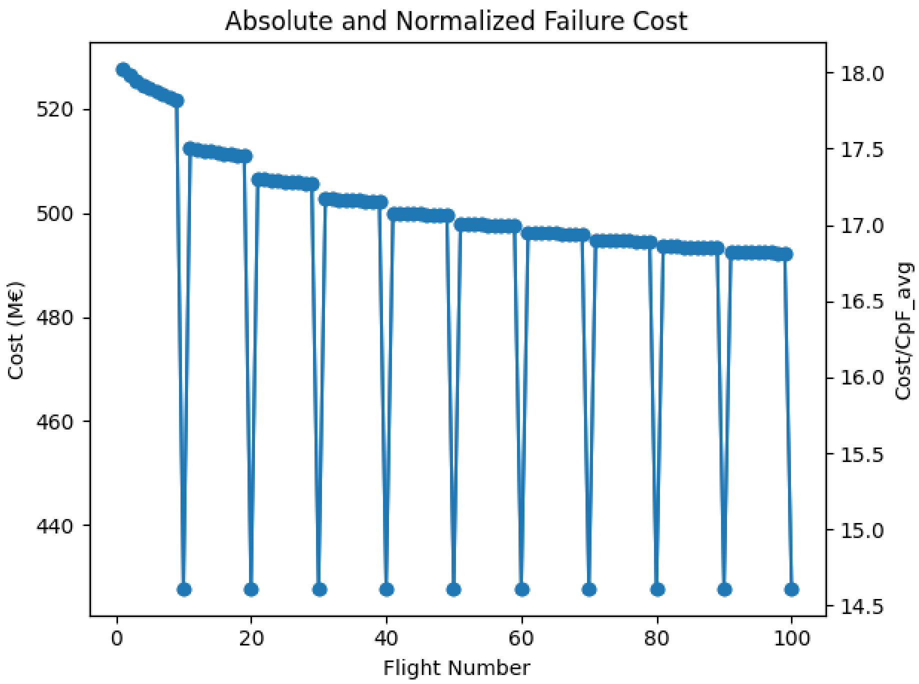

Figure 5.9: Falcon 9 Failure Cost.

5.4. Reliability

Estimating reliability for a commercial launch vehicle like the Falcon 9 is a complicated endeavor. The main reason for this is the lack of transparency inherent to SpaceX being a private company, which makes data-driven investigation very tough to do.

In a development phase, a figure for reliability could be obtained by resorting to two different types of sources: Firstly, data related to different mechanical and electrical systems, present in sources such as MIL-HDBK-217 [93]. The other option would be to use similar systems for an analogous comparison. For instance, in the case of engine reliability prediction, it is common to use the data related to historic systems with similar mass, vacuum thrust, propellant and cycle type [132].

In this case, the model is adapted to deal with a number of assumptions, and the data used is the one pertaining to the operational track record of the Falcon 9 vehicle [45, 46], as well as historical reliability values coming from literature. For this part, only the Full Thrust version of the Falcon 9 is considered. The version currently flying is the block 5 version. Previous versions to the Full Thrust are too different from the block 5 iteration, and so their flight history is disregarded, as it is assumed that redesigns and upgrades have occurred that change too much its reliability.

5.4.1. Engine

As was discussed in previous chapters, the engine reliability metric is one of the most important figures for reliability, as it has historically been the most prominent cause of failure [121, 132]. Since there is not reliability data pertaining to the Merlin engine, nor any information regarding its test and qualification process, apart from some news articles reporting a sporadic failure, or the conclusion of testing, the only information available concerns its operational track record. Even in this, it is complicated to assess the frequency of failure. This happens because the Falcon 9 has engine-out capability. This means that unless the company discloses that a benign failure has occurred, the launcher is capable of accomplishing its mission and in theory safely land. To date, no catastrophic engine failures have been recorded, and the disclosed benign failures are connected to recovery failures, as the stage did not have enough thrust or was incapable of restarting its engines to perform the landing. It is not far-fetched therefore to presume that more benign failures might have happened, that just have not been disclosed, as both the mission and the recovery were accomplished.

In order to derive a value for the engine reliability, only the disclosed failures will be taken into account. It is therefore important to safeguard that this is probably an overestimate of the engine reliability. In the future, when more data is available, it shall be possible to refine this estimate.

From the data available [1], there is a record of 464 successful individual Merlin engine flights, with 6 benign failures. Four engines failed on their first flight, one on its fifth and another on its sixth. As there are more than 2 failures, a Kaplan-Meier estimator can be used to estimate the reliability of the engine. The result is displayed in figure 5.10.

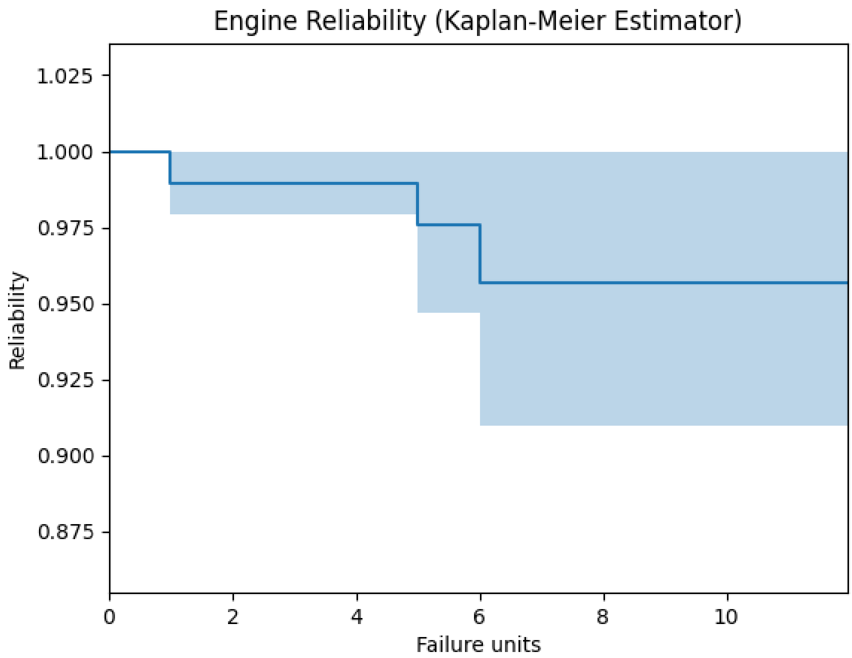

Figure 5.10: Reliability estimate obtained with a Kaplan-Meier estimator (95% confidence bounds)

As can be seen, the confidence bounds are very wide. This happens because there is an extremely high number of right censored data, compared to few failures.

It is assumed that the engine's survival fraction follows a Weibull distribution. This is adequate, as the goal is to model a varying failure rate. By using a Maximum Likelihood Estimator, it is possible to obtain the parameters for the Weibull distribution. The resulting parameters are $\lambda = (182.96)^{-1}$, $\kappa = 1.0707$, and the survivor function is illustrated in figure 5.11. The uncertainty telegraphed by the very wide 95% confidence bounds is very high. In studies with zero or few failure data, usually a 50% confidence bound is used [139, 144, 146]. However, as indicated in Nelson [144], this should serve as a useful warning about the uncertainty of the failure lacking data, rather than discouragement, and so the 95% limit is preferred.

[1]Pertaining to the Full Thrust version of Falcon 9, of April 20th 2022

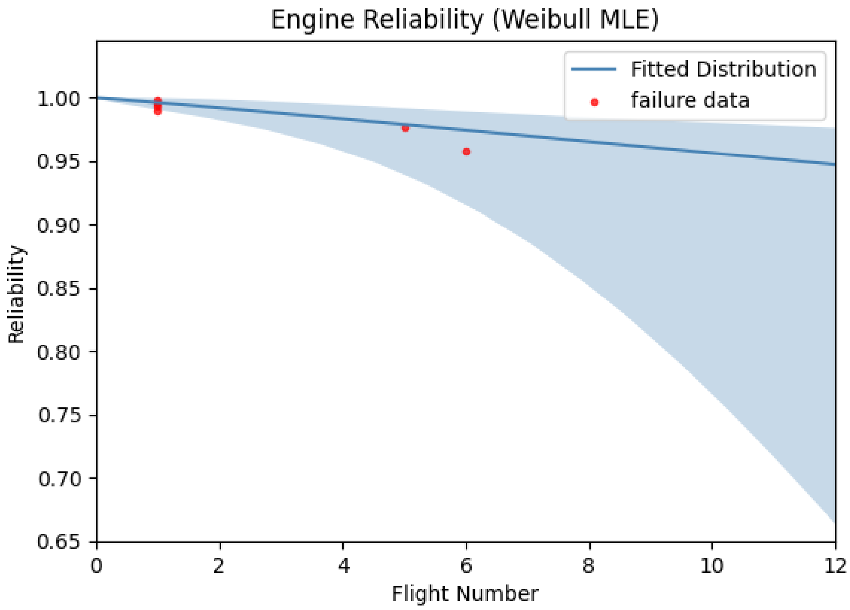

Figure 5.11: Engine reliability estimate obtained with a MLE, assuming a Weibull distribution (95% confidence).

Comparing the Kaplan-Meier estimate with the parametric estimate, illustrated in figure 5.12 it seems that the behavior of the Weibull distribution follows the expected survivor function obtained with non-parametric methods. With this data set being heavily right-censored, no other models are attempted, and the most simple model is used. In future studies with more detailed and more extensive data, goodness-of-fit of several different models can be compared in order to achieve the best one.

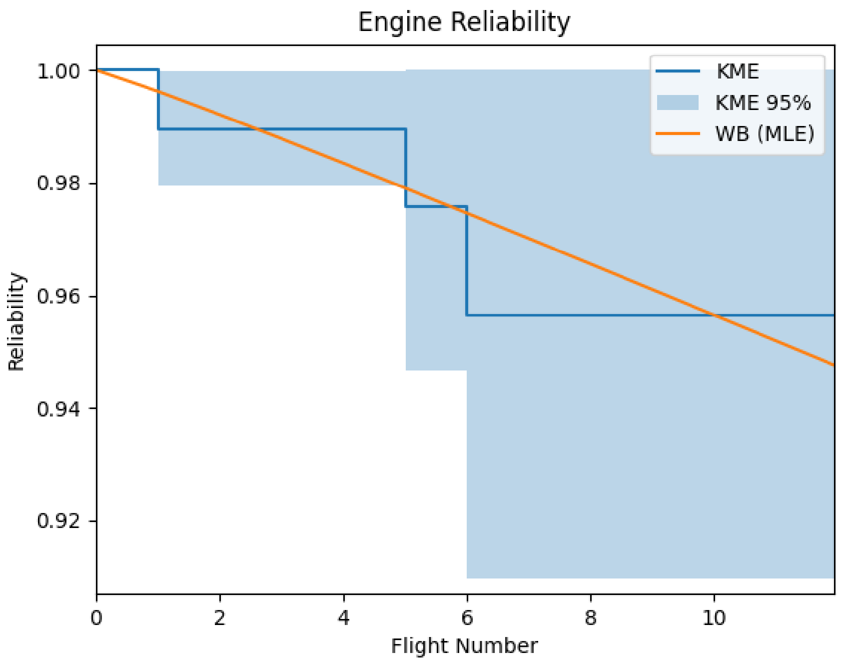

Figure 5.12: Comparison of the non-parametric estimate with KME and the parametric estimate with MLE assuming Weibull distribution

Over a life-cycle of 10 flights, the average reliability for an engine with these characteristics is 0.9766. Comparing to the reliability estimate of the Rocketdyne J-2 engine [132] of 0.988, this signifies an error of 1.15%.

5.4.2. Propulsion Engine-Out Capability

As has been mentioned already, the engine-out capability is a very important feature of the Falcon 9 vehicle. Not only it results in an increase in reliability, especially compared to the 9 engine configuration without engine-out reliability, but it also potentially plays a role in concealing the real failure track record of the Merlin engine.

Applying equation 3.26, the reliability of the first stage propulsion is calculated. The common cause failure percentage and the catastrophic failure percentage are both assumed to be 10%. The average of the propulsion system reliability over a life-cycle of 10 flights is 0.9627, which presents a difference of 1.44% to the S-II main propulsion system (which had a reliability of 0.9768) and 2.05% to the S-IVB main propulsion subsystem (0.9828) [132].

5.4.3. Remaining Subsystems

Since there is no testing or operation data that provide information about systems other than propulsion, and knowing that historically they aren't as relevant to failure as propulsion, the remaining subsystems are modeled with an exponential distribution, assuming constant failure rates found in literature, from the Saturn V data and the projected reliability for the Ares V vehicle [132]. The selected mean flights between failure are summarized in table 5.12.

Subsystems	MFBF
Structures	909
Avionics	1000
Power	6250
Fairing	3096
Separation	25000
Other	100000

Table 5.12: Mean flights between failure based on Saturn V and Ares V.

5.4.4. Recovery

In the Falcon 9 Full Thrust flight history [45, 46], there have been 114 successful landing and only 8 failures. Of these, 5 pertain to the aforementioned engine problems (insufficient fuel or thrust, restart problems, unexpected shutdowns), and the remaining three were due to TVC failure, incorrect wind prediction and grid fin failure.

This failure data was used to generate a non-parametric estimate with a Kaplan-Meier estimator, and is represented in figure 5.13:

In the figure, it is possible to see that the reliability estimate is highly biased downward. This happens because there is a large amount of right-censored data for the first landing (13 out of 30 right censored data points) and second landing (8 out of 30 right censored data points). Furthermore, five out of eight failures happened on the first landing attempt. In addition to that, in the versions up to Block 4, the maximum number of landing attempts for the same booster was only two, with all three failures happening on the first landing. For that reason, only the Block 5 version shall be considered, as it has boosters with up to 12 successful landings and a higher variety of failures.

Assuming that the recovery survival is ruled by an exponential distribution, a Maximum Likelihood Estimate is performed with the failure and right-censored data pertaining to the Falcon 9 FT Block 5, and the result is plotted with the failure data in 5.14, and alongside the non-parametric estimate in 5.15.

The derived constant failure rate λ is 0.055, which translates to an expected 18 mean flights between failure.

Once again, this is a rather imprecise result. This is a feature of firstly the lack of available data at subsystem level, and secondly, of the SpaceX fleet inconsistency. As there is a preference on the client side (especially for military missions) to fly on new boosters, there is a lot of boosters that have landed few times, which results in a very high proportion of right censored data. Coupled with a relatively few amount of failures, that mostly happened on the first landing attempts ($\leq$6), this leads to a result that is heavily biased on the side of unreliability especially on the earlier flights. This could hint at a the presence of failures related to immaturity of the system, however, since only the block 5 is being considered, and there is a disproportionate amount of boosters flying less than 6 times, a simpler distribution with constant failure rate is preferred. In future work, when more boosters have flown, with a greater variety of right censored data and potentially more failures, it would be worth considering a non-constant failure rate, or better yet, estimating at subsystem-level.

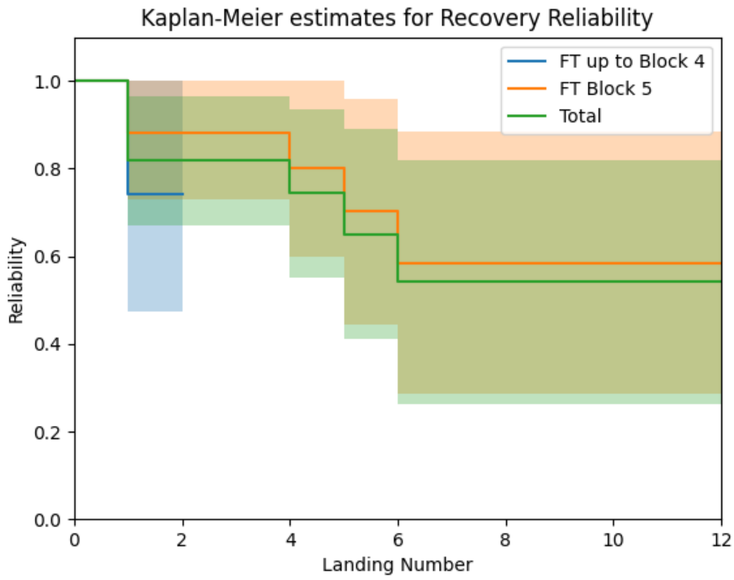

Figure 5.13: Recovery reliability estimate obtained with a Kaplan-Meier estimator (95% confidence bounds), for Falcon 9 Full Thrust up to Block 4, Block 5, and both combined

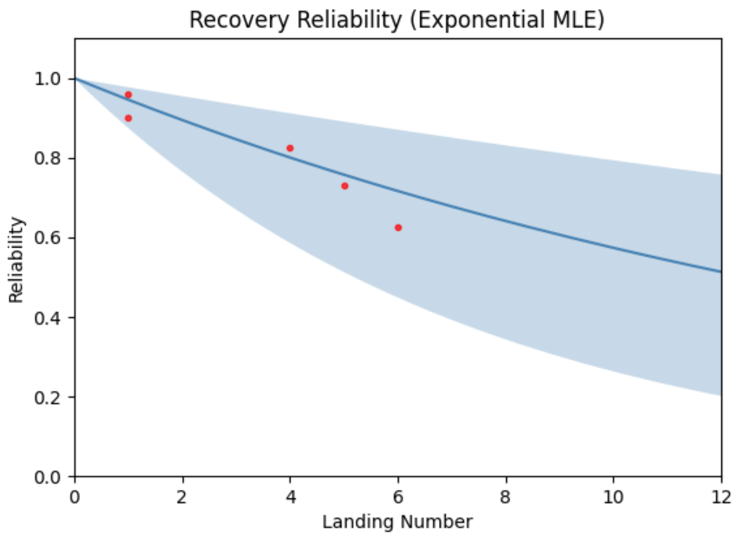

Figure 5.14: Recovery reliability estimate obtained with a MLE, assuming an Exponential distribution (95% confidence).

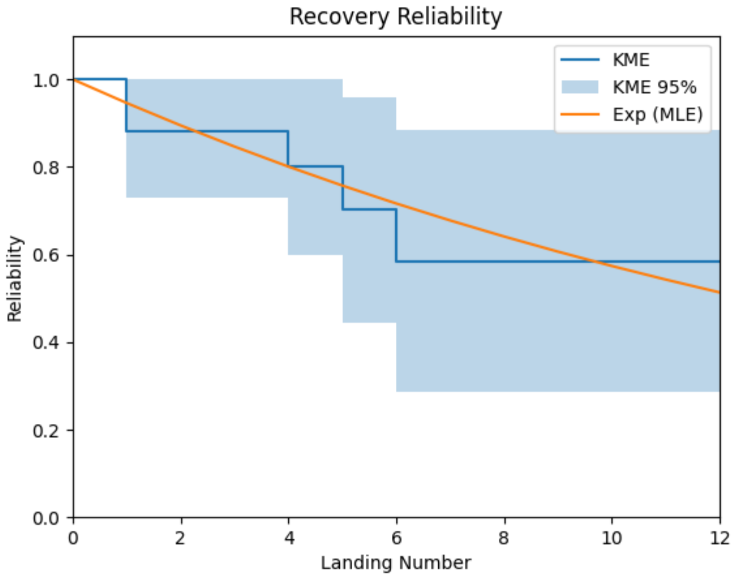

Figure 5.15: Comparison of the non-parametric recovery estimate with KME and the parametric estimate with MLE assuming exponential distribution.

5.4.5. Reliability Analysis

According to the methodology described in section 3.4.2, the reliability calculation at system level is achieved with fault-tree analysis (FTA). The top event is considered the loss of mission (LOM), which occurs in case there is a catastrophic failure in either of the stages, the fairing or in the separation events. This is depicted in figure 5.16.

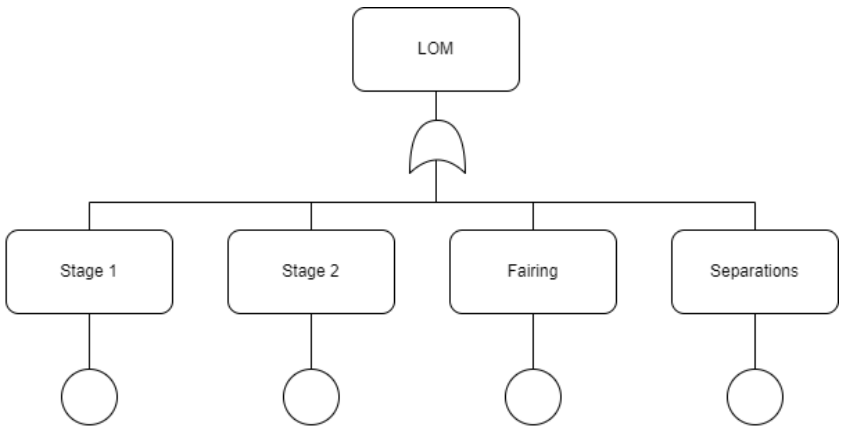

Figure 5.16: Fault Tree Analysis at stage level.

The stage is further broken down into its subsystems. The main subsystems considered are power, avionics (both with redundancy), propulsion, structures, and finally, a generalization of the remainder contributors to failure. This is represented in figure 5.17

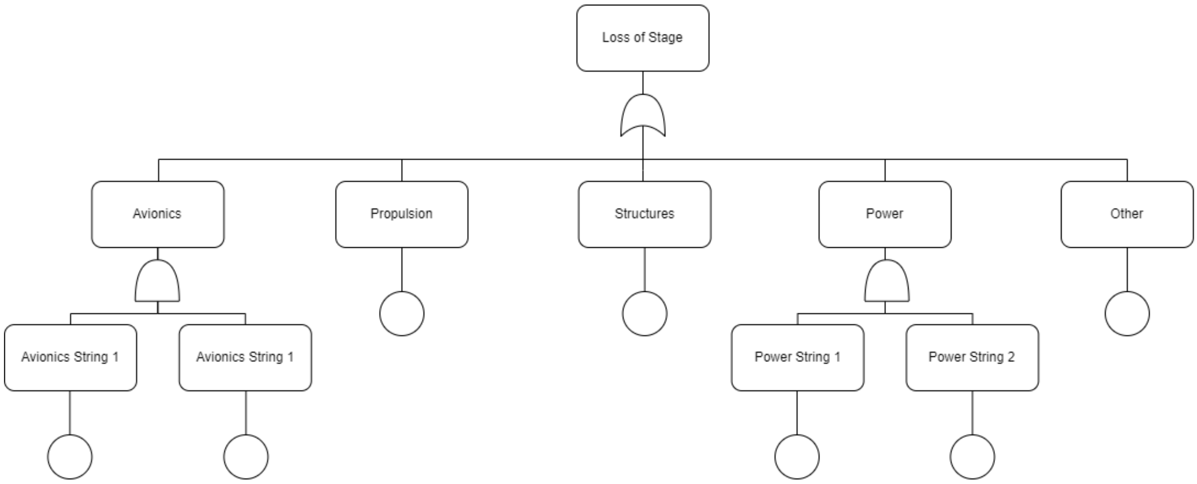

Figure 5.17: Fault Tree Analysis at subsystem level.

Finally, special attention is given to the propulsion subsystem, as a main contributor to failure. The engine-out capability is modeled with a 2-out-of-9 gate, and the model includes common-cause failure and catastrophic failure. This is illustrated in the fault tree in of figure 5.18

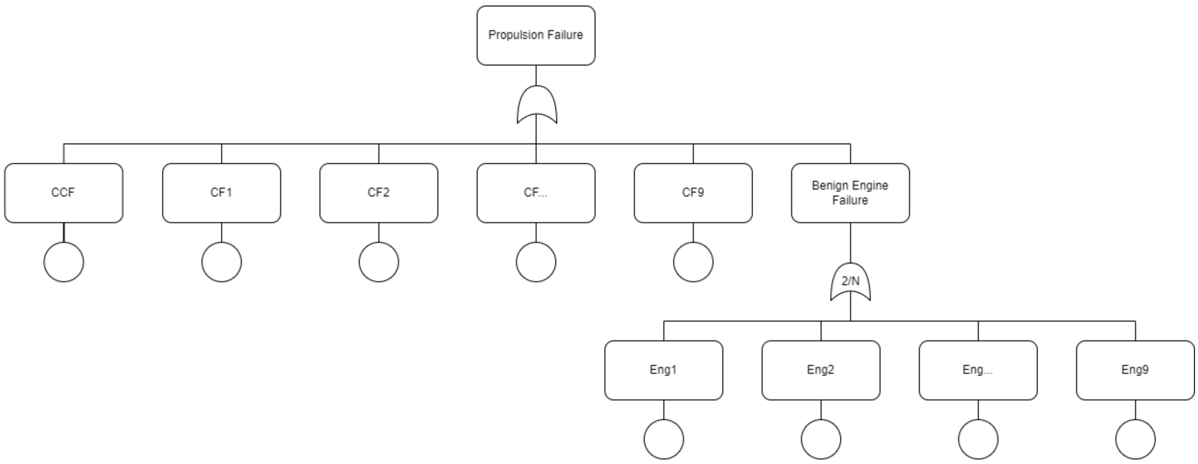

Figure 5.18: Fault Tree Analysis at engine level.

5.4.6. Reliability Growth

Finally, following the evidence demonstrated in section 3.5.1, a reliability growth factor of 0.4 is applied, which means assuming that there is a program dedicated to failure elimination.

5.4.7. Results

After estimating the subsystem reliability, analyzing the systems reliability with FTA, and applying growth, the results are propagated to a program life time of 100 flights, with the reusable hardware having a life-cycle of 10 flights before being expended and re-manufactured. The results obtained are summarized in table 5.13, and illustrated in figure 5.19, including mission and landing reliabilities, and the probability of recovering the first stage.

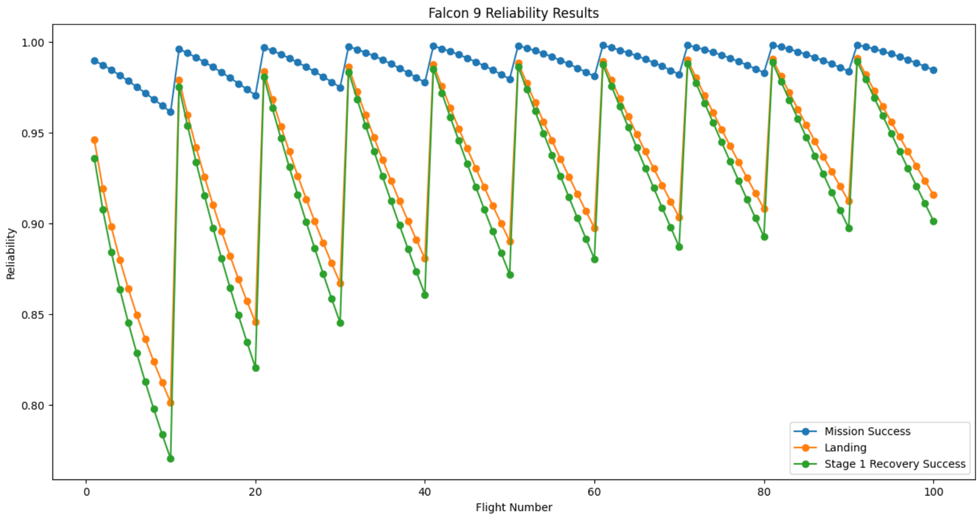

Figure 5.19: Falcon 9 reliability results, including mission and landing reliabilities, and the probability of recovering the first stage

Element	Average Reliability
Mission Success	0.9881
Landing Success	0.9298
First Stage Recovery	0.9190

Table 5.13: Reliability results summary.

As there are no sources on the Falcon 9 reliability, these numbers are compared against some estimates found in literature. The findings are summarized in table 5.14.

Element	Estimated Reliability	Present method error	Source
Falcon I Expectation	0.9716	1.70%	Futron Corp. 2004 [87]
Falcon V Expectation	0.9738	1.47%	Futron Corp. 2004 [87]
Falcon 9 Mission Estimate	0.9717	1.69%	Martino 2010 [149]
Falcon 9 Bayesian Estimate	0.9380	5.34%	Everett and Dezfuli 2016 [108]

Table 5.14: Reliability results verification

Observing the graph represented in figure 5.19, one can see how much less reliable the landing is compared to the mission. This is consistent with the expectations, as the Falcon 9 Full Thrust version has never had a single failure, (excluding an explosion on the launch pad caused by problems in fueling), compared to several landing failures.

In addition to this, both the mission and landing reliability display a very pronounced degradation effect throughout the life-cycle of the reusable hardware. Some ways to counteract this effect are to use better quality parts (more reliable throughout the operation interval), to increase the effectiveness of the refurbishment process (less degradation), and to proceed to the replacement of some critical parts (comparable to fully repairable systems). All of these strategies imply however an added cost.

5.5. Expected Cost of Failure and Value

Having now access to the reliability figures for both the mission and landing, it is possible to combine them with the estimated failure cost, and come to the expected cost of failure, following the methodology outlined in 4.8.

The results throughout the lifetime are represented in figure 5.20, for the different insurance cases. The average costs of failure considered throughout 100 flights for the different cases considered are summarized in table 5.15.

Case	$C_{F_{avg}}$ (M€)	$C_{F_{avg}}/CpF_{avg}$ (%)
Simple	9.4	32.1%
RFG	9.7	33%
INS	7.2	24.4%
RFG+INS	7.4	25.4%

Table 5.15: Expected Cost of Failure result summary.

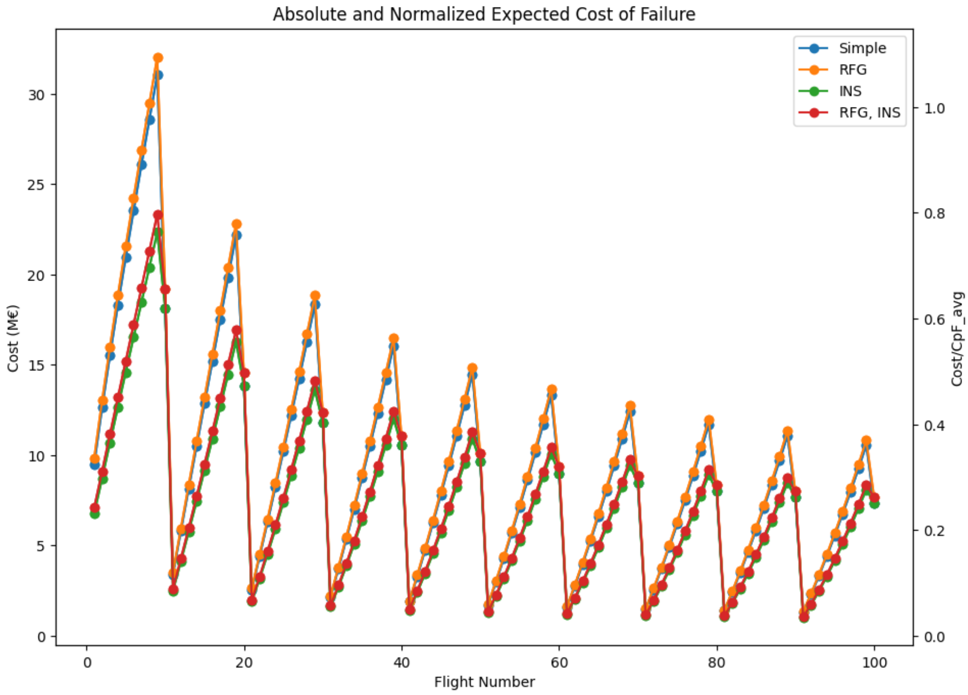

Figure 5.20: Falcon 9 absolute (left) and normalized (right) Expected Cost of Failure.

From the figure, it can be seen that as the reliability increases, throughout the whole life-time of the launcher existence, and as the cost per flight also decreases, the expected cost of failure tends to in general decrease with the number of boosters flown.

Particularly in the flights of the first booster, the expected cost of failure can reach values close average cost per flight, depending on decisions regarding insurance, for the last flights in the booster's useful life.

The values of the expected cost of failure throughout a booster's life-time tend to reduce to around 0-40% of the average cost per flight, depending on the age of the booster.

Looking at the different scenarios considered, it can be seen that the alternative that results in a higher expected cost of failure is when there is a re-flight guarantee agreed between the launch provider and the satellite operator. This is to be expected, as the responsibility to ensure a new launch falls with the launcher company. It is important to remember that this computation does not yet include the premium paid by the satellite operator for this type of insurance. Following the same trend, the result that yields the lowest expected cost of failure is when the launch vehicle reusable parts are insured.

Taking the probability of success and profit into account, the Expected Value can be computed according to equation 4.55.

The expected value was calculated in two modes: The first considers a constant profit margin of 48%, which corresponds to the average profit obtained. The second mode considers SpaceX's pricing scheme, which is a fixed price of $62M for launch on a new vehicle, and $50M for a launch on a reused one. As the cost per flight decreases with the number of launches, this situation sees the profit per flight increase.

The results are plotted in the graph presented in figure 5.21. The results have been normalized with the average profit. The average results of the expected value are summarized in table 5.16

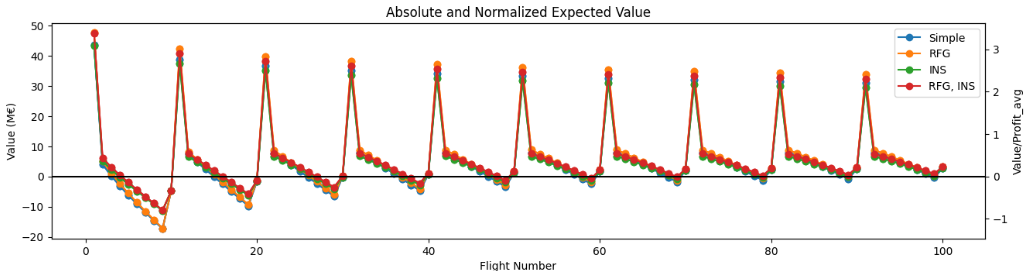

(a) Expected Value with constant profit margin.

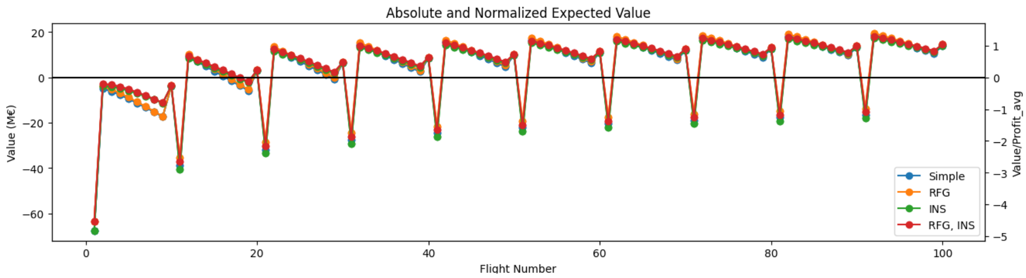

(b) Expected Value with constant price.

Figure 5.21: Absolute and normalized expected value.

Case	EV_{avg} (M€)	$EV_{avg}/(Profit)_{avg}$ (%)
Simple	4.6	32.8%
RFG	5.5	39%
INS	5.0	35.7%
RFG+INS	5.8	41.4%
Maximum Value	6.2	44.3%

Table 5.16: Expected Value result summary.

The average expected value for the Falcon 9, throughout its lifetime and not counting for any type of insurance other than liability, is 4.6 M€, which corresponds to 33% of the average profit obtained in a normal cost study, not including failure. If the action that maximizes the expected value for each flight is taken, then it is possible to increase this figure to 6.2 M€, which corresponds to 44% of the average profit. On the other hand, if only the flights with a positive value were to be flown commercially, maximizing the expected value obtained, then an average expected value of 11.8 M€ is obtained, corresponding to 84% of the average profit.

The first and most striking observation is that not all flights have a positive value. This means that on those flights, the launch provider can be expected to lose money.

In the case where the profit margin is constant, as represented in in 5.21a, the first flights of the life-time of each booster have always a positive expected value. This is due to the fact that this flights is very costly, but also yields a big profit in case of success The aging effect and consequent deterioration in reliability on the reused flights lead to situations where the expected value is negative. Furthermore, as the reliability grows with operational experience, there is a higher tolerance built for the last flights in the life of the booster, to the point where eventually all the flights will have a positive expected value.

If a constant price is considered, as is represented in 5.21b, which more accurately portrays the real situation seen in the industry, the first flight of each booster actually presents a negative value, as it was determined that they are being sold at a loss, according to this model. Contrary to the constant profit margin case, the first booster does not have any flight where the expected value is positive. However, the system itself reaches a point where all the flights using reused boosters yields a positive expected value, while the other model shows that the last flights in the life of a booster are usually not advantageous.

In order to analyze the difference between the different insurance cases, the expected value of the first booster versus the last booster are compared side by side in figure 5.22. The figure represents the constant price situation, however the same trends analyzed hold true for the case of constant profit margin.

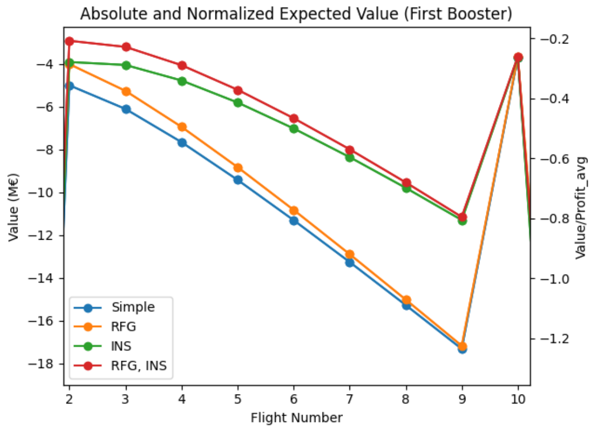

(a) Expected Value of first booster.

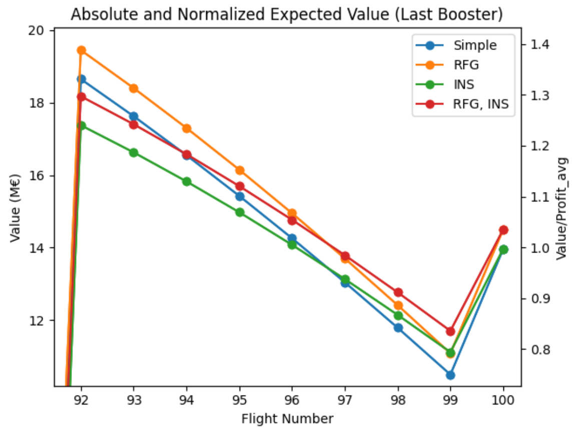

(b) Expected Value of last booster

Figure 5.22: Absolute and normalized expected values of the flights using reused boosters, for the first and last boosters.

For the first flight of each booster, both cases see the insured flight have a lower expected value. This happens because the flight is very costly to insure. The most advantageous case is the RFG, as there is a lesser probability of failure, compared to flying with a reused booster, and the premium is proportional to the cost of the flight, which is higher for a new booster.

The situation differs for the flights when the boosters are reused. For the first booster, it is more advantageous to have the first stage insured throughout its life-cycle. This happens because there is a higher probability of failure, as the system hasn't had time to mature and increase its reliability, and so it is worth to pay the insurance premium as a hedge to mitigate losses. This is followed by the RFG case, as the high premiums paid by the satellite provider compensate the added cost in case of failure. The most beneficial case is the combination of insurance and RFG, and the least is the simple case.

In the last booster however, there is no visible benefit in insuring the booster for the first seven flights. As the reliability degrades, it becomes beneficial to take the insurance, in order to mitigate the more likely losses. Once again, the RFG agreement is beneficial throughout the lifetime of the booster.

This exercise shows how reliability is a very important parameter to account for in a cost and economic study. By including the probability of failure and the different outcomes possible, other than a simple cost/price balance, the actual profit that a company can expect (discounting insurance) is actually closer to a third of the nominal profit yielded from the difference between price and cost over the life-time of the vehicle. On the other hand, if it is assumed that the flights with negative value are not commercialized, with the risks being covered by a public agency for instance, then the expected value of the flights that are indeed commercial average out to a figure closer to 85% of the average profit, with only 15% being lost to failure, when considering the fixed price case. With a constant profit margin, the expected value in these circumstances would be lower, at 8.4M€, or 60% of the average profit. This difference is due to the fact that in the different modes, there is a different number of flights with negative expected value, and by eliminating these points, the average obtained is not the same. In this situation, it be would be necessary to increase the profit margin to roughly 75%, from 48%, in order to achieve an average expected value that is equal to the average profit, when not accounting for reliability and the cost of failure. This would result in an average price per flight of 51.3M€, which represents a price increase of 18.2% from the originally considered 43.4M€ ($51.2M). This price increase of 7.9M€, or 27% of the average cost per flight, although not being a cost inherent to the vehicle development, manufacturing and operation, is a factor that can be added to the price paid by the customer in order to reflect the expected cost of failure.

There are a number of reasons that can explain the large amount of flights with a negative expected value.

Firstly, regarding the failure cost model, it is assumed that on average a medium-effort modification is necessary after failure. However, it would not be unreasonable to assume that as time progresses, major and medium modifications would become less frequent, and as a result only minor modifications would be done long-term, or even no modifications. This is an assumption that is common in FMECA analysis [90].

Secondly, in the sphere of reliability, it was already noted how the lack of testing data leads to results that have a high uncertainty. On top of that, given the limited amount of engine and landing failures, and their prevalence in the first few flights in the lifetime of a booster, it is possible that the launcher lifetime reliability is being underestimated with this method. If this is true, then it would be reasonable to expect that the initial reliability of the launcher is higher than the one forecast, and that the reliability growth effect would sooner bring the flights to a favorable expected value range. On that note, the reliability growth factor selected is higher than that obtained from historical data coming from expendable launcher experience. If it is the case that this assumption is wrong, then one could expect that the reliability growth would be slower, and more flights would yield a negative expected value.

Overall, the results obtained here correspond to a reasonable depiction of reality, if one is to look to the operational history of Falcon 9. For the Full Thrust versions up to Block 4, no booster was launched more than 2 times, having either been destroyed upon landing, expended, or retired after their first or second flight. This era spanned 27 boosters and 39 successful flights. Only in the current Block 5 iterations can we find boosters being flown a multitude of times, with the maximum having been 12 launches. The observation that it took multiple boosters with 1-2 flights until a nominal state of multiple flights per booster was reached supports the thesis that maturity and reliability growth is required in order to reach an operational state where multiple flights per booster have a positive expected value. This is also coherent with the expendable launch vehicle experience at ESA, that sees a multitude of reliability enhancing techniques in the early phase of operations, such as redundancy, which is suppressed long-term as the launch vehicle proves its reliability [148].

Finally, the expected value computation is necessary yet insufficient measure for decision-making in a commercial environment. The utility of the launch requires further study into the attitude of the company regarding risk, as a risk averse company will require a more demanding reliability requirement (less risk) and an assured payoff (positive expected value), whereas a more risk proverse company will tolerate lower reliability (more uncertainty) for a relatively higher value, or even a negative value for a comparably greater payoff.

5.5.1. Comparison to Expendable Case

Based on the original reusable design, an expendable version is obtained, by shedding the reusable hardware and the costs related to reusability. It is expected that this version will have a higher reliability, but higher cost per flight. Both these factors will have a direct impact on the expected cost of failure.

The main differences implemented in order to obtain the expendable version and its direct effects on cost are summarized as follows:

- Removal of the recovery hardware and removal of 30% of original landing mass in propellant weight [165];
  - DEV = 355 M€;
  - MAN_{avg} = 45.5 M€
- Changes to the failure cost model:
  - Length of investigation: 28 days;
  - Length of downtime: 120 days;
- Average reliability of 0.9895.

The model is applied for a yearly launch rate of 10, common for expendable launch vehicles, and 20, closer to the higher launch rates flown by the Falcon 9. Two different profit margins are also studied, the first equal to 8%, following ESA guidelines, and the other the same as the average profit obtained for Falcon 9, 48%.

The main findings are summarized in table 5.17.

LpY	10	10	20	20
Profit (%)	8	48	8	48
CpF (M€)	58.3	58.3	65.5	65.5
PpF (M€)	63.0	86.3	70.7	96.9
C_F (M€)	3.2	5	1.5	2.3
C_F/CpF (%)	5.5	8.6	2.3	3.5
EV (M€)	1.4	23	3.8	29.1
EV/Profit (%)	30.0	82.2	72.5	92.6

Table 5.17: Results for Expected Cost of Failure and Expected Value for the expendable version of Falcon 9.

Firstly, it can be seen as expected that the cost per flight is higher than that of the reusable version. At 20 flights per year, the cost per flight is 124% more expensive. Considering the same average profit as the reusable version of 48%, the price per flight reflects the same difference.

In terms of the expected cost of failure, these represent a much lower fraction of the cost per flight, all lower than 10%. This is a result of the higher reliability of the expendable version, compared to the reusable one, in spite of its higher costs.

Looking at the expected value results, the proposal with the highest expected value is the one that mirrors the situation obtained for the reusable version of Falcon 9, with the yearly launch rate of 20 flights and the profit margin of 48%. Although this result is much higher than the one obtained for the reusable version, being 633% higher, it also results in the much higher price per flight mentioned before. This would mean that operating this vehicle, at the same launch rate and same profit margin of its reusable counterpart, the launch provider could expect gain more money with the expendable version. The caveat however is the higher price per flight. This model does not take into account the effect of price per flight on the launch rate, which in effect would represent the preference of the customer on the launcher. It can be expected that an increase of price would lead to a decrease in the launch rate, as the launcher becomes less competitive. In that case, keeping the profit margin constant and considering a launch rate of 10 flights per year, it can be seen that the expected value of the expendable version drops by 10.4%, being still higher than the reusable version, but also being almost exactly twice as expensive. If the profit margin is reduced to 8%, making the price per flight of the expendable solution only 45.3% more expensive than the reusable one, then the expected value drops to 1.4M€, which is 70% lower than the value of the reusable version.

This exercise shows that the higher reliability and lower expected cost of failure of an equivalent expendable launch vehicle appears to have a higher expected value, being therefore more profitable and preferential to a reusable launch vehicle, given the same conditions financial conditions (same profit margin) and same level of demand in the market (expressed by the same launch rate). However, given that the expendable solution isn't as competitive from a cost standpoint, it can be expected that the launch rate would be lower, and that the profit margin wouldn't be as high, in order to lower the price. These actions lead to a decreasing expected value, and as the price of the expendable version approximates the figures of the reusable version, the expected value becomes a small fraction of the expected value obtained for the reusable configuration.

5.5.2. Sensitivity Study

Launch Rate

The launch rate is an important parameter, as it directly influences the cost of operations, the amortization of the development cost, and the failure cost. The results are summarized in table 5.18.

Parameter	Variation	Parameter Value	CpF avg (M€)		CF avg (M€)		EV avg (M€)	
			Value	Error (%)	Value	Error (%)	Value	Error (%)
LpY	-50%	10	37.7	28.7	6.5	-30.9	-0.9	-120.0
	+50%	30	26.1	-10.9	13	38.3	4.2	-8.7

Table 5.18: Launch rate sensitivity study.

It can be seen as expected that a higher launch rate leads to a lower cost per flight. This has been derived in literature many times before [99, 129, 132]. However, due to the prominent effect of the launch rate on the failure cost model, its increase leads to a higher expected cost of failure. This comes from the losses stemming from downtime, as there is no revenue from launches being sold, and the accumulating facilities costs. A curious observation is that in both cases surveyed, the expected value obtained is lower than when a yearly launch rate of 20 flights is assumed. This hints at the existence of an optimum yearly launch rate, that balances the decreasing cost per flight with the increasing losses due to failure.

This sensitivity analysis results in two important points: firstly, the observation that the facilities cost reference used stems primarily from TRANSCOST and ESA sources, which are calibrated for lower launch rates. This could result in a disproportionate cost of failure when considering much higher launch rates as those currently seen on the market (such as the +30 yearly flights observed by SpaceX). This is something that can be refined in a future study. Secondly, the fact that higher launch rates require leaner facility and failure investigation solutions, which that minimize the losses due to downtime. This can be achieved with streamlined and efficient investigation practices.

Finally, the fact that the increasing launch rate eventually leads to a negative expected value due to the increasing launch rate, despite the ever decreasing cost per flight showing the importance of including reliability and failure when analyzing feasibility of a launcher project from a cost perspective, as past research in the area has generally recommended to strive for the highest launch rate possible [99, 127–129, 133, 158, 185].

Engine Reliability

The case where the engine reliability on the first flight is equal to 0.999 is investigated. The Weibull distribution parameters are kept constant, so this constitutes a proportional increase, before the effect of reliability growth. The results are summarized in table 5.19. This exemplifies how a small reliability difference of only 0.3% has a positive impact in the estimated value of two orders of magnitude. This opens the door to trade-offs in the development phase, as will be exemplified in chapter 6.

Parameter	Variation	Parameter Value	CpF avg (M€)		CF avg (M€)		EV avg (M€)	
			Value	Error (%)	Value	Error (%)	Value	Error (%)
R_{eng_i}	0.30%	0.999	29.3	0	8.8	-6.4	5.2	13

Table 5.19: Initial engine reliability sensitivity study.

Landing Reliability

The same exercise is repeated for the landing reliability. An increase to 0.999 in the initial landing reliability corresponds to a 5.6% increase, which has a similar impact to the increase in engine reliability. The results are presented in table 5.20.

Parameter	Variation	Parameter Value	CpF avg (M€)		CF avg (M€)		EV avg (M€)	
			Value	Error (%)	Value	Error (%)	Value	Error (%)
$R_{landing_i}$	+5.6%	0.999	29.3	0	8.6	-8.5	5.4	17.4

Table 5.20: Initial landing reliability sensitivity study.

Catastrophic Failure Percentage

The catastrophic failure percentage, which directly impacts the first stage reliability, as it is an important parameter in the engine-out capability model is tested at 0.2 and 0.3, values found in literature for expendable launch vehicles [121, 132]. The results, gathered in table 5.21, illustrate how these proportionally high variations have a more moderate impact in the final expected value balance, but proving however the value of reducing this figure to the minimum possible in order to maximize the expected value.

Parameter	Variation	Parameter Value	CpF avg (M€)		CF avg (M€)		EV avg (M€)	
			Value	Error (%)	Value	Error (%)	Value	Error (%)
CF	+100%	0.2	29.3	0	11.4	21.3	2.6	- 43.5
	+200%	0.3	29.3	0	13.4	42.6	0.6	-87.0

Table 5.21: Catastrophic failure percentage sensitivity study.

Reliability Growth Factor

The reliability growth factor is an important parameter in this study, as the viability of reusing a booster multiple times is predicated on the assumption that operational maturity leads to an increase in reliability, thereby mitigating the aging effects, as was demonstrated. By using a growth factor of 0.2 (-50%), which has been derived from historical data, and denotes that ”corrective action is taken for important failure modes” [132], a negative expected value is obtained, effectively limiting the amount of times that are financially viable to reuse a booster. The results are summarized in table 5.22. This study demonstrates the importance of a running a program dedicated to failure elimination, in order to maximize the expected value.

Parameter	Variation	Parameter Value	CpF avg (M€)		CF avg (M€)		EV avg (M€)	
			Value	Error (%)	Value	Error (%)	Value	Error (%)
α	-50.00%	0.2	29.3	0	18.1	92.6	-4.0	-187.0

Table 5.22: Reliability growth factor sensitivity study.

Profit Margin

The constant profit margin, originally taken at 48%, is varied between -50% and +50%. It is shown that, by raising the profit margin, although the cost of failure is raised, the expected value obtained from it outweighs that loss. On the other hand, by reducing the profit margin by half, the project becomes inviable, with a negative average expected value. This shows that there is a certain average profit margin that needs to be applied, in order to make the commercial launch vehicle viable, given a certain reliability, and that in this case, it exceeds the standard 8% dictated by ESA. The results can be consulted in table 5.23

Parameter	Variation	Parameter Value	CpF avg (M€)		CF avg (M€)		EV avg (M€)	
			Value	Error (%)	Value	Error (%)	Value	Error (%)
Profit Margin	-50%	0.24	29.3	0	8.7	-7.4	-1.7	-137.0
	+50%	0.72	29.3	0	10.1	7.4	11.0	139.1

Table 5.23: Profit margin sensitivity study.

5.6. Lessons Learned for the Future European Reusable Launch Vehicles

Given the results obtained from the application of the cost and reliability model to the Falcon 9 launcher, and the subsequent failure and expected value analysis, a set of discoveries are uncovered, which can serve as lessons learned and recommendations for the future European reusable launch vehicle programs, such as Themis [12].

- As was already stated in SOLSTICE [99], minimizing the variation in TRL needed for the development of the launcher equipment leads to cost savings. For that reason, use of mature technology from the different ESA programmes will lead to cost savings in the development phase. This includes not only the heritage from the Ariane family of launch vehicles, but also the new reusability technologies deployed in the IXV and SpaceRider.
- As was found in research, the cost estimates of Falcon 9 by NASA using NAFCOM went from around $1.7B to a $440M. The main factors that contributed to this reduction of 74% in the estimate included technical corrections, an updated mass breakdown, and insight gained from a visit to the SpaceX facilities. As the estimate reduction occurred mostly at Design, Development, Test and Evaluation level (-80%) than with the test units itself (+9%), which are directly dependent on the T1, then it is reasonable to assume that the brunt of the change is carried by the assumptions regarding the development program, rather than the ones concerned with the mass breakdown. This indicates that the development process conducted at SpaceX is considerably different and slimmer than the one assumed by the NASA cost experts, when the first estimates were produced. This is compounded by the fact that these changes resulted from a visit to the SpaceX facilities, which provided the experts with insight into the procedures. What this indicates is that there is a large space to find cost saving measures in the usual way governmental agencies such as NASA have lead their development efforts. An assessment of the way development is carried out in ESA projects is therefore recommended, in order to isolate inefficiencies and approximate these procedures to the way they are carried out in the private sector.
- In terms of manufacturing, the flights that include a new booster have a disproportionate high cost of manufacturing, when compared to the flights using a reused booster. These costs can be brought down by resorting to commonality with the components present in the upper stage: as the upper stage is produced every flight, having common hardware between the upper stage and the first stage, as is the case with the engine in the Falcon 9 example, makes it so that the gains derived in the upper stage production from the learning curve every flight will also be applicable to the first stage equipment. This will be further explored in chapter 6.
- Regarding operations, the indirect operation costs are constant, and don't apply to public contracts (such as those provided by NASA or the DoD) or launches of satellites owned by the launch provider, as is the case of SpaceX launching Starlink satellites. The direct operating costs have the highest share in operating costs, and benefit from a learning effect. It can be expected that the first few flights in the life-time of the vehicle will be the most expensive ones in the operations domain. Regarding retrieval costs, the launch rate was identified as the most important driver. Furthermore, the key to lowering the vehicle recurring costs lies with lowering the refurbishment costs. The Falcon 9 concentrates these costs in the engines. Technological developments at engine level appear to also be going in the direction of minimizing the refurbishment intensity and

frequency, as the Methalox combustion is expected to leave less soot, therefore eliminating the need for recurrent and extensive cleaning operations.

- As for pricing, it was found that the flights with brand new hardware could be commercially sold either at cost, or even at a loss. This would indicate an expectation to have the flights with reused boosters compensate for the rest of the flights. This effectively means that the profit for the launch provider is coming from the latter type of flights. As the first Full Thrust versions before block 5 flew no more than two times, this indicates that most of the revenue came from the public clients, whose flights were reported as being sold at a mark-up [22, 26, 40, 42, 62]. This would in turn help build confidence in the system, and lead to reliability growth. This pricing strategy also allows the Falcon 9 to beat the market price of the competition, as the most expensive flights have slim (or negative) profit margins, while the lower costs achieved by reusing the system allow higher profit margins.

- Continuing on pricing, the findings showed that if the profit margin falls below a certain point, the enterprise assumes a negative value. Since the typical ESA profit margin is counted at 8%, for this specific launcher and its characteristics, this would result in an average negative value. As ESA is a public entity, which should prefer value neutral propositions, then two scenarios occur: Either the profit margin is adjusted in order to reflect a positive value; or the lifetime of the reusable hardware is limited to only include flights where the result of the value computation is positive, therefore not using the equipment to the best of its durability

- The failure model developed and its application allows to look at the sources of cost and plan to mitigate them. Whereas some costs are inevitable, such as the flight substitution, and some are unpredictable and out of the launch provider's control, such as climbing insurance rates, the costs related to the downtime and investigations can be optimized. The best way to do this is by minimizing the amount of downtime, and therefore limiting the amount of flights missed, and the investigation costs themselves. This in turn can be achieved by having in place a systematic approach to the failure investigation, with systematic procedures studied and implemented. On top of that, streamlining the time it would take to develop the necessary modifications and implementing them would also contribute to minimizing the failure cost. On the other hand, cost of downtime can also be minimized by minimizing the recurrent costs pertaining to keeping the facilities and staff operational during downtime.

- Although previous studies have shown that the cost benefit of reusable launch vehicles is directly related to their launch rate [99, 127–129, 133, 158, 185], this study shows how high launch rates amplify the cost of failure. For that reason, the previously indicated solutions to minimize downtime become even more important. Furthermore, the manufacturing and operation cost reduction and contrasting increasing cost of failure resulting from increasing the launch rate suggests the existence of an optimum launch rate per year. As the cost of failure stemming form downtime is minimized, this optimum shifts in the direction of an increased launch rate, enabling higher cost savings in manufacturing and operations.

- The results coming from the expected value model allow the launch vehicle provider to obtain a measure of how much capital is expected to be gained (or lost) upon a certain flight. Companies can use this information, according to their risk tolerance philosophy to make informed decisions regarding their flight policy. For a risk neutral entity, as ESA would be classified, only flights with a positive expected value should be flown. The results showed that early in the lifetime of the reusable launch vehicle, it is not viable to fly a stage multiple times, as the high costs and reliability degradation lead to a negative expected value. Although these flights could still be considered "useful", for instance for a risk favoring (proverse) company, it would be advisable to focus on dealing with reliability and allowing it to grow with a program dedicated to failure elimination. For this, as was already referred, a dedicated program supplying the launcher with the necessary flight rate to do these improvements and gain experience and rank in the learning curve is desirable.

6 Multi-engine application

The trend in the current launch-vehicle industry seems to lean towards multi-engine configurations to propel their first stages, as some of the most prominent commercial launch vehicle companies have been implementing them in their flagship launch vehicles. SpaceX's highly successful Falcon 9 launcher relies on nine Merlin engines in its first stage [4]. Rocket Lab's Electron also depends on 9 engines (Rutherford) [3]. The reusable vehicles in development of these companies also foresees a multi-engine configuration in the first stage: the Starship's first stage, the Super Heavy, can house up to 33 Raptor engines [68], which is reminiscent of the soviet N1 launcher, with 30 engines in its first stage [47]. The latter launcher was plagued with accidents, never having a successful flight, and its program was canceled in 1974. The Neutron's first stage houses 7 engines [48].

The next launcher in the Ariane family, Ariane 6, will only count on one Vulcain 2.1 engine, together with two or four P120C strap-on solid boosters [19].

In this chapter, a use-case of the cost and reliability models is demonstrated, concerning the assessment of a single-engine versus multi-engine first stage configuration. A similar exercise has been previously worked out by Koelle [129]. In this study, it was demonstrated how an increase in the number of engines lead to a decrease in the total life-cycle cost of a launch vehicle. In this chapter, this exercise is revisited, using Ariane 6 and its engines as an example, and this time including reliability as a variable. By considering an engine-out scenario, the thrust of each engine in a certain configuration is necessarily higher than the one obtained by simply dividing the original thrust by the number of engines, as is done in TRANSCOST. Furthermore, by making the reliability of the engine cluster a fixed parameter, additional qualification is deployed in the cases where the reliability of the cluster is lower than that of the single-engine configuration. The total life-cycle cost of the engine system isn't therefore only a product of the number of engines, as it is in the example worked out in TRANSCOST. Furthermore, additional variations are considered: a reusability scenario is analyzed, where some assumptions are taken in order to simulate the life-cycle of the same launcher with recovery and refurbishment of its first stage, in order to see if the final recommendation regarding the first stage engine configuration changes. Furthermore, an analysis on the advantages of engine commonality between the first and upper stages is done, in both the expendable and reusable configurations, in order to assess the impact of this measure on cost.

Lastly, this model does not have the capability to address some issues directly, such as expected increases in mass, which would require detailed performance, propulsion and trajectory models in order to converge to a result. For that reason, it is recommended that this model be considered for an optimization tool, and the caveats mentioned will be here investigated in a sensitivity study.

6.1. Method

In this section, the publicly available data about the Ariane 6 launcher and the Vulcain 2.1 engine that will be used in this study is gathered. Some of these will be further processed, using the method and equations developed for this research project. Finally, a description of the methodology for estimating the cost and reliability of the different configurations is outlined.

6.1.1. Available Data

For this exercise, the Vulcain 2.1 engine will be taken as an example.

The Vulcain 2.1 is the engine that will power the first stage of the Ariane 6 launcher, in a single-engine configuration. It is a liquid cryogenic engine, currently being developed by Snecma in cooperation with ESA. It is a similar engine to its predecessor Vulcain 2, achieving a similar thrust level, with technological simplifications [78].

Using publicly available data [19, 32, 78, 166], the relevant engine and mission characteristics are obtained:

- Mass: $M_0 = 2150$ kg;
- Thrust (vaccum): $T_0 = 1371$ kN;
- Burn time: 600 s.
- Payload Mass:
  - A62: 10,500 kg to LEO, 4500 kg to GTO;
  - A64: 20,600 kg to LEO, 11,500 kg to GTO.

Information about the qualification program for the Vulcain 2.1 engine is also available, and contains detailed data [166]:

- 3 test campaigns;
- 25 tests;
- 60 ignitions;
- Cumulative testing time: 14465 s.

This program seems to have a relatively number of tests, when compared to the one held for the upper stage engine, Vinci, which was reported as having more than 100 tests [24]. This could be due to the fact that the Vulcain 2.1 engine has design heritage coming from Vulcain 2, therefore needing a shorter qualification process in order to achieve the required reliability targets. This will be further explored in the next section and in the sensitivity analysis present in section 6.2.7.

6.1.2. Data Processing

Firstly, from the testing data available, it is possible to infer the reliability to which the Vulcain 2.1 engine is being tested to. Applying equation 3.22 at a confidence level of 90%, the reliability achieved through testing would be $R_0 = 0.9467$. This seems to be a low value for reliability, especially when comparing to historical references [129, 132]. This could be explained by the fact that this rocket engine is an iteration over the original Vulcain 2, itself based on the first versions in the Vulcain family of engines. This could signify that there is a heritage which brings benefits in both cost and reliability [121, 153]. This heritage is accounted for in the development costs in SOLSTICE, as it would demand a lower ΔTRL, however in reliability this is harder to quantify. Between Vulcain and Vulcain 2, a less extensive test program is observed [153]. In the method, the reliability of 0.95 shall be used, with a Δ TRL of zero being assumed. In the sensitivity study of section 6.2.7, a case where the starting reliability is the same as that of Vulcain 2 engine is investigated (R=0.995 [151]).

From the mass data given, it is possible to get a breakdown of the costs that are influenced by this component. These were obtained by simply following the methodology and recommendations for ELV systems encompassed in SOLSTICE [99]. As it is assumed that the engine possesses heritage coming from the previous iterations Vulcain and Vulcain 2, a Δ TRL of 0 is assumed. This is ultimately not a driving decision in this study, as an increase in the Δ TRL only increases the Engineering Function of the Project Office costs, which is a component of the development cost and is felt across board regardless of the number of engines considered in the configuration.

A period of 100 flights is considered, as that is around the number of lifetime flights of the Ariane 4 and Ariane 5 launchers, and at a flight rate of 10 flights per year (which is more than the current launch rate of Ariane 5 and Vega combined), it is equivalent to 10 years of operations, which is the recommended amortization period for the development costs. The results obtained are listed below, with the manufacturing and operating costs averaged over the 100 flights:

- $T1 = 15.040$ M€;
- $DEV = 103.290$ M€;
- $MAN_{avg} = 8.744$ M€;
- $OPS_{avg} = 0.213$ M€.

6.1.3. Methodology

The methodology described in this chapter is summarized in figure 6.1. It looks to firstly derive the main characteristics of possible multi-engine configurations that are capable of achieving the mission profiles desired for the Ariane 6 launcher and are therefore "equivalent" to the projected single engine configuration. After that, the costs relative to these different configurations are estimated, taking into account certain reliability targets, and an optimal configuration is found, minimizing cost.

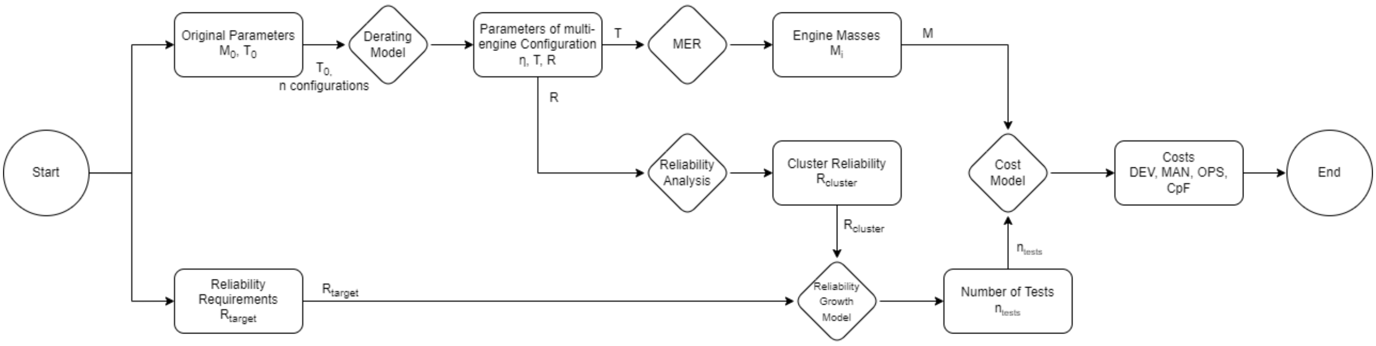

Figure 6.1: Multi-engine problem solving methodology.

In order to extrapolate the figures obtained for the single-engine configuration to the multi-engine configurations, it is necessary to account for the evolution of the engine characteristics, as their number increases. It is assumed that the stage features engine-out capability. This means that when an engine fails, the rest of the engines that incorporate the cluster are able to increase their thrust. The result is that the total thrust of the impaired cluster matches the one achieved in nominal operating conditions. The engines need therefore to be rated to a thrust level that is higher than the nominal one. The thrust of each engine in an n-engine configuration shall be $\alpha_n \cdot T_0$, where T_0 is the nominal thrust of one engine, and α_n is the derating level correspondent to the n-engine configuration.

The derating constant is obtained from equation 6.1. If an engine is derated at level 0.5, it means that the engine is designed to supply 50% more thrust than the non-derated version, in case this is needed.

$$\alpha_n = \frac{1}{n-1} \tag{6.1}$$

The mass of the engines in each configuration is a fundamental parameter for the cost estimation model. The mass per engine can be obtained using a mass estimating relationship (MER), relating thrust to mass [105, 193, 194]. For a liquid, turbo-pump fed, cryogenic engine, equation 6.2 relates mass (in kg) to thrust (in kN):

$$M = 0.006 T^{0.858} p_c^{0.117} (A_e/A_t)^{0.034} \tag{6.2}$$

The range of applicability and relative standard error for this equation could however not be found in the reference from which it was borrowed [105]. The design parameters corresponding to the chamber pressure p_c and area ratio A_e/A_t are assumed to be constant and independent from the number of the engines considered in each configuration. The factor in the equation depending on these parameters are therefore deduced from the known engine data and not altered.

Although this MER provides a mass estimate corresponding to the required thrust, this does not result in a complete thrust/mass model, as an optimization process would be necessary in order to calculate the thrust necessary to account for the extra engine mass and not suffer a penalty in payload capacity. Because of this, in this study, a payload mass penalty is inserted, equal to the added mass from the engines, and it is recommended that in the future this model is integrated in an optimization tool that includes complete mass and performance modules, such as the FRT [165].

As was discussed in chapter 3, the uprating of the engines to a thrust level that is above the nominal value, and their further operation at a derated level, leads to an increase in reliability. This can be obtained using equation 3.28 and assuming a constant failure rate. The reliability of the cluster is obtained considering an engine-out model. The suggested catastrophic failure percentage is between 10-30% [121, 132]. This means that out of the failures that occur at engine level, for a particular engine, a certain fraction is catastrophic, which necessarily result in loss of the launcher. The rest consist of failures which can be mitigated by shutting down the affected engine and operating the remaining engines at a higher thrust, in order to still accomplish the mission. A lower catastrophic failure percentage translates into a higher reliability.

Following this, a reliability target is determined for the engine cluster. The engine reliability that is needed, for the different configurations, in order to achieve the cluster target, can be derived by taking the inverse function of equation 3.26. The number of tests necessary to reach this engine reliability can be calculated using Duane's reliability growth model, introduced by equation 3.21 and considering a factor of 0.4, which signals that there is a program dedicated to failure elimination [132]. The cost involved in this further testing effort can be obtained using equation 4.10, the CER derived from ESA data relating the test program to its cost.

The cost calculation follows from the model already described in chapter 4. A detailed breakdown is examined in order to assess the main contributors to cost. Furthermore, a reusable case is considered. This is done by adding extra non-recurrent costs (related to the recovery hardware), and retrieval cost, and including the refurbishment costs.

For this case, a propulsive vertical landing at the launch site was chosen, as it is the only solution currently in use in the launcher market, and the return to launch site strategy is more cost advantageous than the downrange landing one. A performance study is recommended and necessary in order to determine the most advantageous strategy considering the mission profile and the launcher characteristics. Specifically, if the horizontal velocities involved in the recovery flight allow a return to launch-site landing [172]. As the recovery mode is not a driving factor in this study, the most straightforward and convenient strategy is assumed. The added costs that need to be accounted for are the development of the landing equipment, the retrieval costs, and the refurbishment costs of the engines and the landing gear. The landing gear costs are estimated following the methodology described in the previous chapter, and considering that the ratio between the engine cluster mass and the landing legs mass is constant and equal to that of the Falcon 9 results obtained in the previous chapter. There are additional cost savings that are not analyzed in this comparison, namely the value that is extracted from reusing the remaining hardware of the first stage, rather than manufacturing it for every flight. Since there is a lack of information concerning the mass breakdown of the Ariane 6 stage, and that this study focuses mainly on the engines, these cost savings will not be considered here.

6.2. Implementation, Results and Discussion

In this section, the implementation of the methodology described in the previous section is analyzed, detailing important assumptions and recommendations for future work.

The results obtained are presented, with an important focus on their effects on the final cost per flight of the different configurations.

6.2.1. Derating level, Thrust and Mass

The evolution of the derating level of each engine depending on the number of engines in a certain configuration is illustrated in figure 6.2.

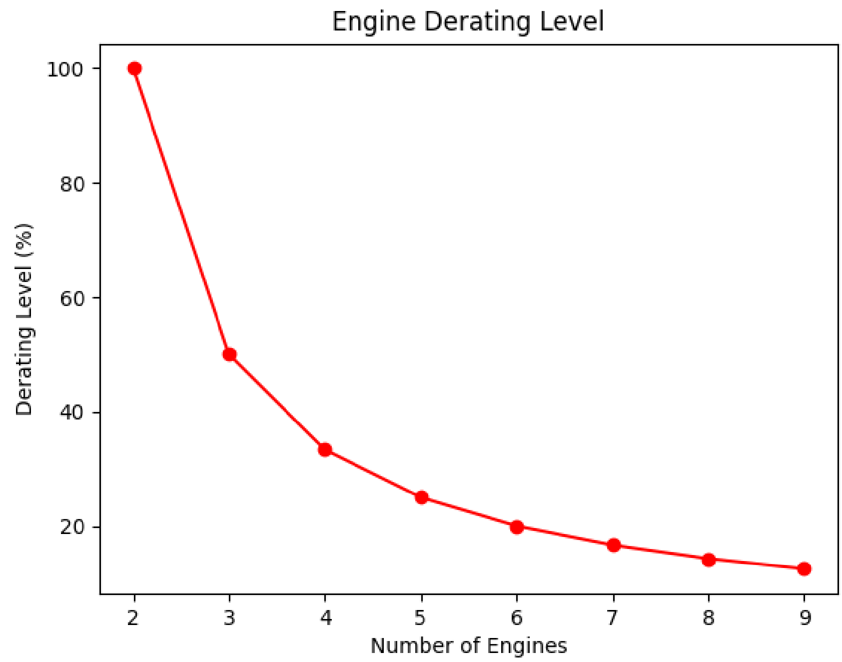

Figure 6.2: Derating level corresponding to the number of engines in each configuration.

Following from the engine-out assumption, the thrust to which each engine needs to be rated in order to be able to accomplish the mission in case of failure of one engine in the cluster is depicted in figure 6.3.

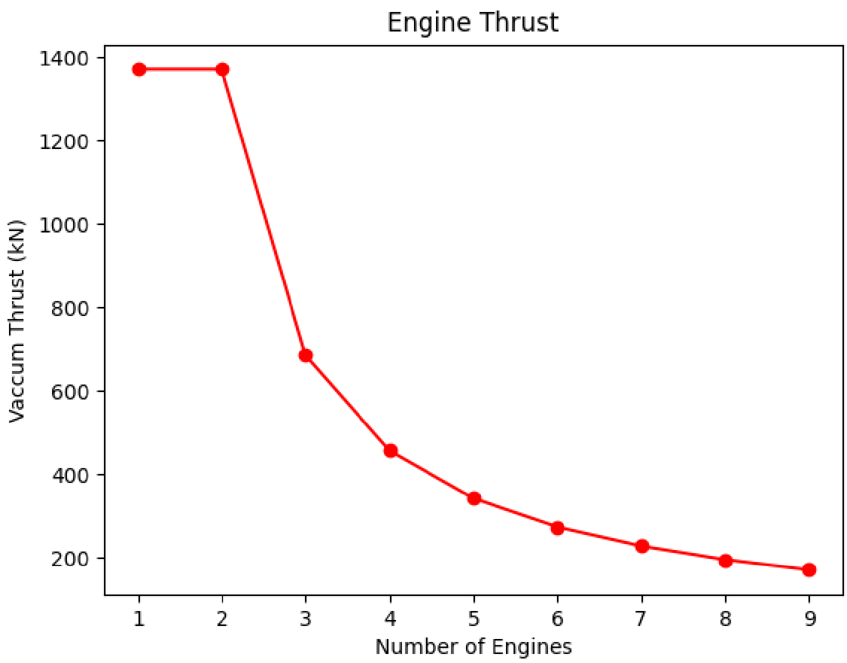

Figure 6.3: Thrust to which each engine is rated in order to enable engine-out capability.

The first observation is that the case where n=2 illustrates a simple redundancy. This makes sense, since if one of the engines fails, the other one needs to be able to achieve the thrust required to complete the mission with one single engine. As was mentioned in the methodology, no thrust adjustment is done to account for the extra mass of the second engine, as the increase in mass of the cluster is accounted for as a penalty in the payload mass.

This means that the two engines installed in the first stage are equal to the one present in the single-engine configuration. This will have cost and reliability related consequences: The theoretical first unit cost and therefore the manufacturing cost per engine shall be the same in both configurations. At reliability level, there will be no increase due to derating, but there will be a reliability increase in the cluster due to redundancy.

Secondly, as the number of engines in the configuration increases, the slope of the derating level tends to zero. Since the thrust, and therefore mass depend on the derating level, this means that as the number of engines increase, the characteristics of the engine tend to stabilize, and they become more homogeneous.

A caveat in this method is that the impact of the cluster geometry in the thrust uprating is not considered, and only the total thrust is contemplated. This means that the position of the engine where a failure occurs does not matter, and the model cares only that the resultant thrust of the launcher when its engines are working above their nominal rating after a failure is equal to the resultant thrust in a nominal state. Furthermore, possible mass differences coming from shortening the aft skirt with the increasing number of engines, or differing thrust frames, are not considered.

A final observation is related to the fact that, as the number of engines increases in the configuration, the vacuum thrust of a single engine becomes very close to the thrust of the second stage engine, Vinci [91], which has a vacuum thrust of 180 kN. This opens up a potential for cost savings, as by using the same engine in both the first and second stage, it only becomes necessary to develop one engine, and there is an increase in the ranking in the learning curve, for the same number of flights.

The resulting mass of a single engine and cluster are depicted in figure 6.4.

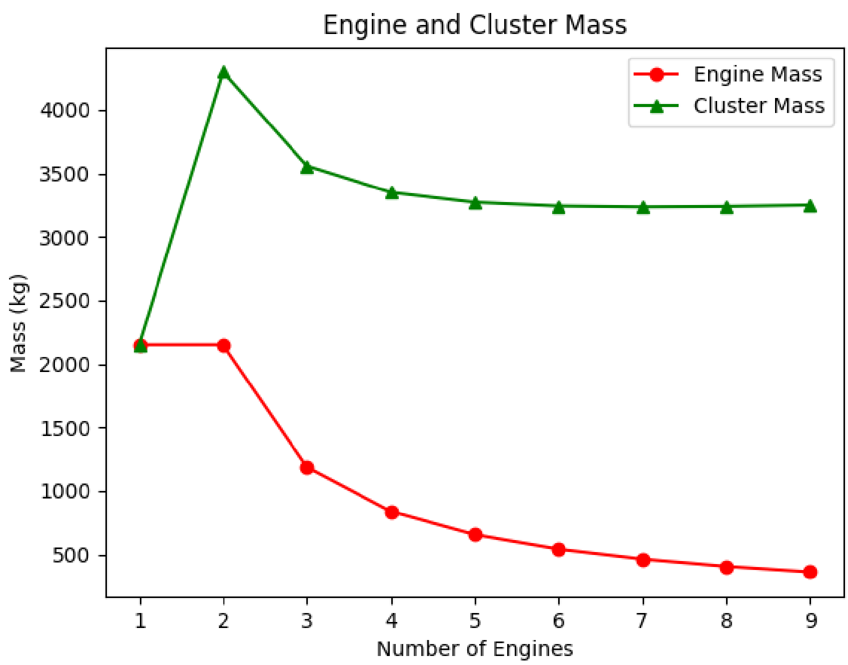

Figure 6.4: Evolution of the mass of a single engine and cluster with the number of engines in the configuration.

The results show that as expected, the mass of a single engine decreases from n=3 until it stabilizes. Regarding the cluster, it is shown that as the number of engines increases, the mass of the cluster of engines first has a steep increase, due it representing a simple redundancy case, and then decreases until an optimum is reached for n=7. From then on, it rises again with a less aggressive slope. This is expected, as the mass per engine stabilizes, but the number of engines in the configuration continues increasing. The slope would be more aggressive, if the payload capacity were to be conserved, as the added structural mass would lead to more propellant being needed. This can be refined in a further study, integrating this model with a performance and propulsion model such as [165].

It is important to note that the mass of a cluster of engines is always greater compared to a single-engine configuration. This will translate in a cost increase, due to loss of payload mass. This payload capacity loss can be directly quantified using the known mass figures for the Ariane 6 vehicle, and is depicted in figure 6.5

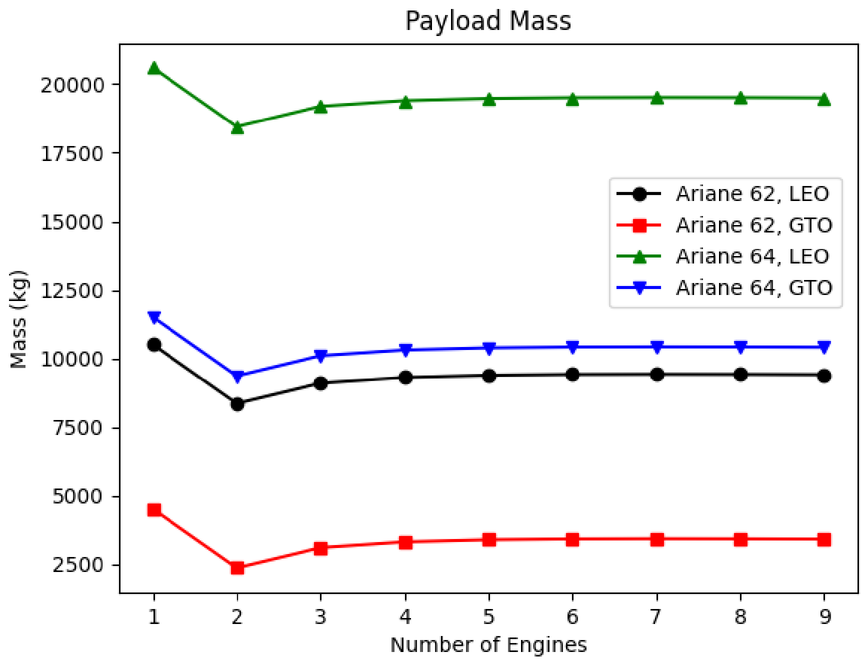

Figure 6.5: Variation of Payload Mass with the number of engines in a given configuration.

It can be seen that the mass penalty quickly stabilizes, being a relevant decrease in terms of relative loss for the case where payload capacity is lowest, namely an Ariane 62 launch to GTO.

The results pertaining to this subsection are summarized in table 6.1.

Number of Engines	1	2	3	4	5	6	7	8	9
Derating Level (%)	-	100	50	33.3	25	20	16.7	14.3	12.5
Thrust (kN)	1,371	1,371	686	457	343	274	229	196	171
Engine Mass (kg)	2,150	2,150	1,186	838	654	540	462	405	361
Cluster Mass (kg)	2,150	4,300	3,559	3,351	3,272	3,242	3,235	3,239	3,250

Table 6.1: Summary of results pertaining to derating, thrust and mass.

6.2.2. Reliability

Using equation 3.28 to calculate the engine reliability increase due to derating, and equation 3.26 to compute the cluster reliability with engine-out capability, the reliability values are obtained and plotted in figure 6.17.

From the figure, the engine reliability sharply increases when n=3, and then the derating effect fades as the number of engines increase. Looking at the reliability of the cluster, there is immediately an increase for the simple redundancy case, which is then decreasing as the number of engines increase. Between 2 and 6 engines, an increase in reliability is observed when compared to the single-engine reliability. After that, the multi-engine cluster reliability becomes less advantageous when compared to using a single engine.

The result of the cluster reliability is sensitive to the catastrophic failure percentage parameter. Literature suggests around 30% for expendable vehicles [132], and less for reusable vehicles [121].

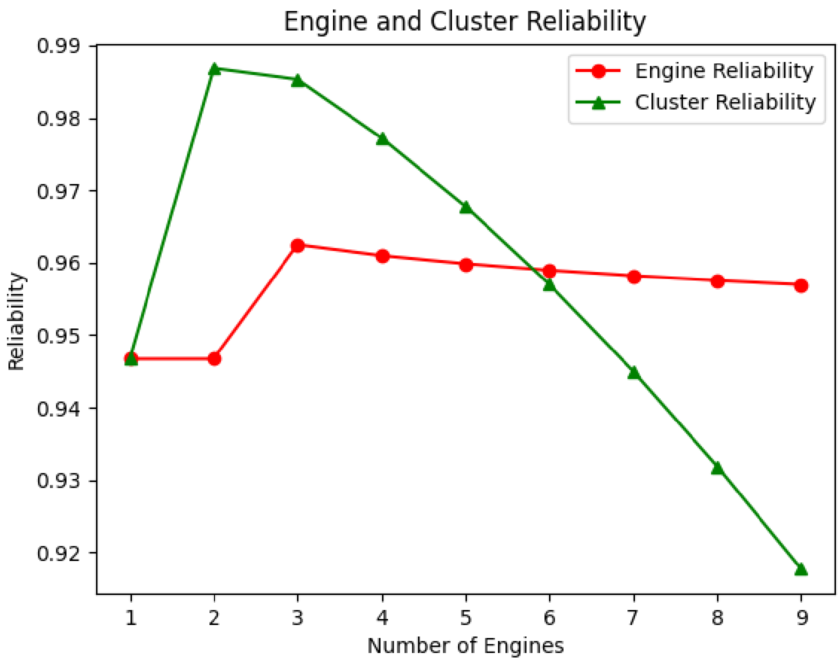

Figure 6.6: Single-engine reliability resulting from derating and engine cluster with engine-out capability reliability, as a function of the number of engines in configuration (CF=0.1).

This fact is illustrated in figure 6.7.

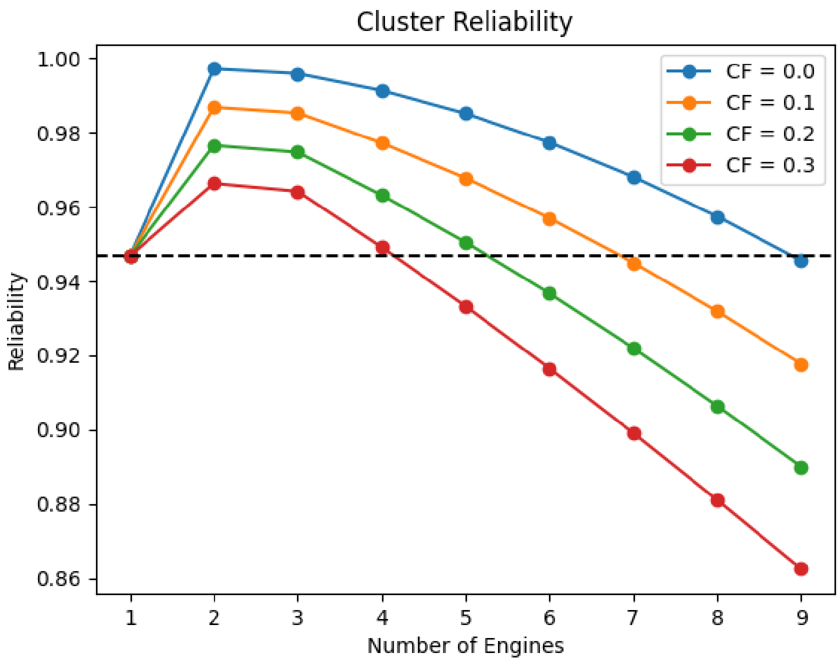

Figure 6.7: Sensitivity of cluster reliability to the catastrophic failure percentage (CF).

This sensitivity analysis shows that depending on the catastrophic failure percentage considered, different multi-engine configurations are deemed more advantageous from a reliability standpoint than the single-engine solution. At CF = 0.3, up until 4 engines there is a reliability improvement. At CF =

0.2, the n = 4 configuration is made viable, and at CF = 0.1, there are improvements registered until n = 6.

Taking into account a target for the reliability of the cluster, the number of cycles necessary to achieve values of reliability in the interval [0.990; 0.995] are investigated. Using the Duane reliability growth model, it is possible to calculate the number of ignitions necessary to achieve these targets. This is depicted in figure 6.8.

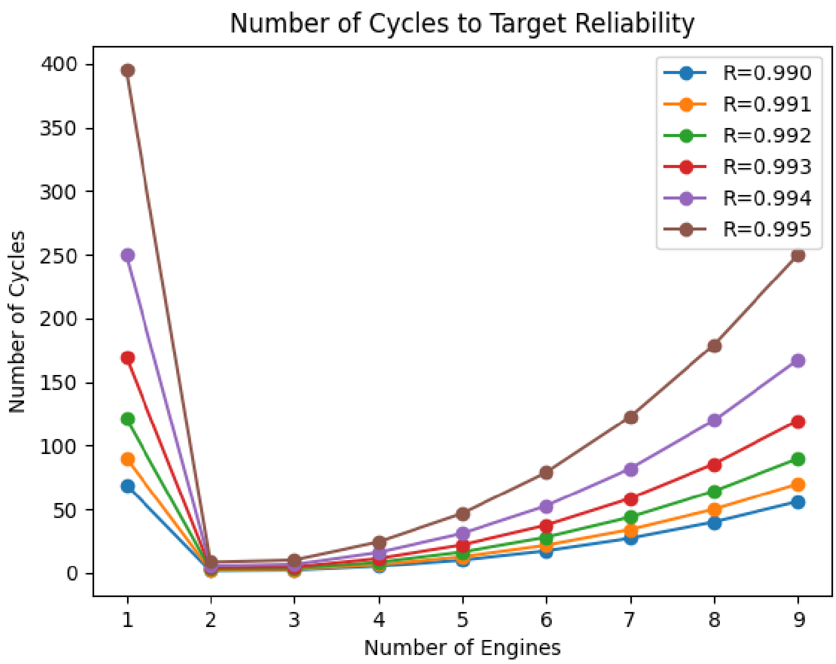

Figure 6.8: Number of cycles needed to reach reliability target

The cost involved in this further testing effort can be obtained using the CER derived from ESA data 4.10. This will be done in the next section.

The results obtained in this subsection are summarized in table 6.2.

Number of Engines	1	2	3	4	5	6	7	8	9
Engine Reliability	0.9467	0.9467	0.9625	0.9610	0.9598	0.9589	0.9582	0.9575	0.9570
Cluster Reliability	0.9467	0.9868	0.9853	0.9772	0.9678	0.9570	0.9450	0.9319	0.9177
Ignitions (R=0.990)	69	2	2	5	10	18	27	40	56
Ignitions (R=0.991)	90	2	3	7	13	22	34	50	70
Ignitions (R=0.992)	121	3	4	9	17	28	44	64	90
Ignitions (R=0.993)	170	4	5	12	22	38	59	86	120
Ignitions (R=0.994)	250	6	7	16	31	53	82	120	167
Ignitions (R=0.995)	395	9	10	25	47	79	123	179	250

Table 6.2: Summary of results pertaining to reliability (CF=0.1).

6.2.3. Cost

The cost estimates can be obtained trivially following the methodology described in chapter 4, once the mass has been calculated. These results are summarized in table 6.3. Figure 6.9 demonstrates the evolution of the theoretical first unit cost for the engine cluster for the different configurations, figure 6.10 illustrates the development costs that follow from that.

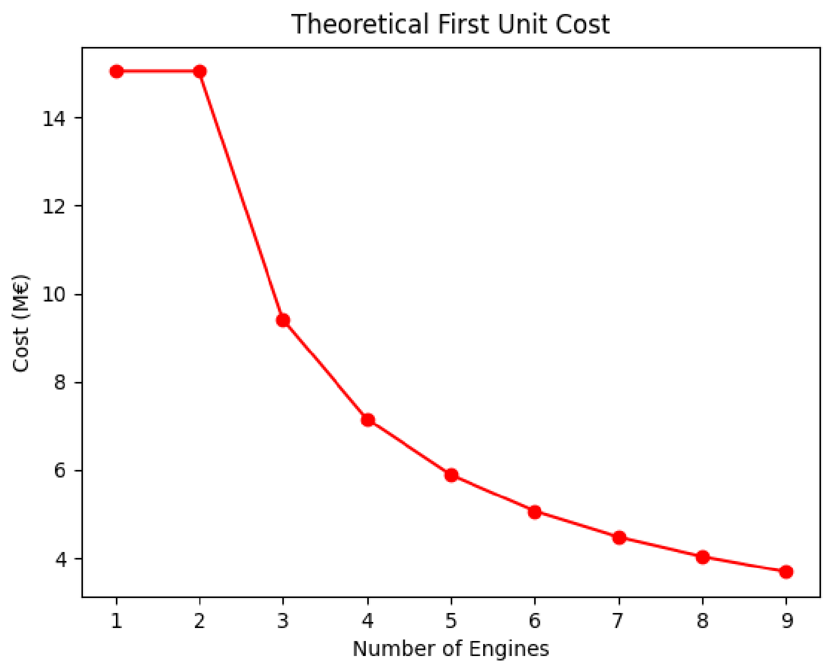

Figure 6.9: Engine theoretical first unit costs variation with the number of engines.

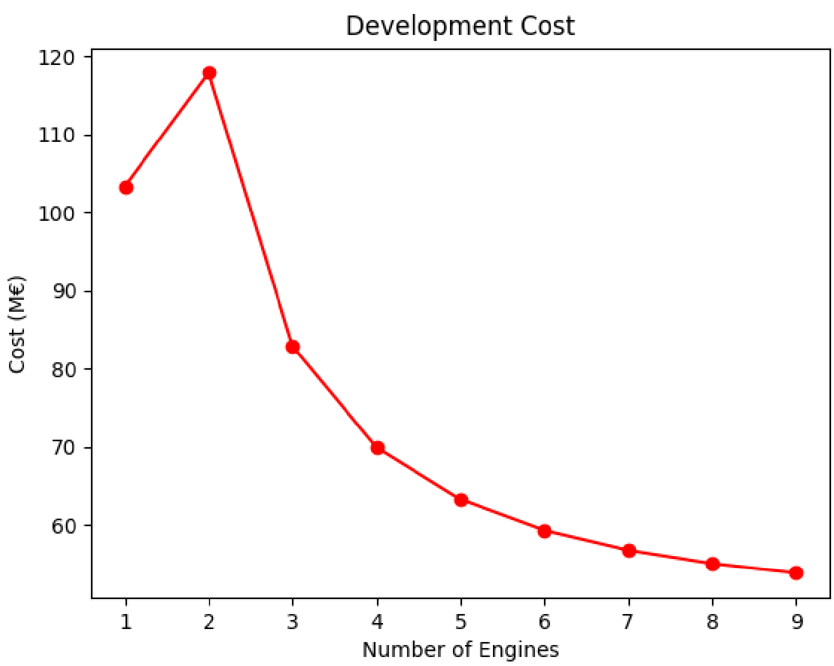

Figure 6.10: Engine development costs depending on the number of engines.

The engine T1 is naturally decreasing as the mass of the engines decrease. Once again, the T1 of the engines for the 1 and 2 engine configuration is the same, as the engine is effectively the same with simple redundancy.

The engine development costs depend on the T1 but also on the quantity of development models. As expected by analyzing equation 4.9, even though the n=1 and n=2 corresponding engines are the same, the latter is slightly more expensive to develop, as it is necessary to account for the amount of times the system is used in the launcher, which will influence the number of model units that are produced for testing. A multi-engine configuration beyond the simple redundancy case brings a clear benefit in terms of development costs.

The added testing effort needed to increase the reliability can be translated into a cost figure using equation 4.10. It is assumed that a month is made up of four weeks, and that there is a rate of 1.25 tests per week (taken from an internal ESA reference), and each test comprises an average of 2.4 ignitions [166]. Due to the regression obtained, there are negative cost results for these added testing programs that last less than 2 months, therefore lying outside the range of applicability of the CER used. For that reason, these costs are considered to be zero in these cases. In these cases, in order to obtain a more refined result, it would be necessary to either estimate the costs of the particular test campaign, in a bottom-up fashion, for instance according for headcount and the duration of the campaign. Alternatively, determining a minimum operating cost for the test facility would allow to attribute a cost to these shorter campaigns. The results are plotted in figure 6.11.

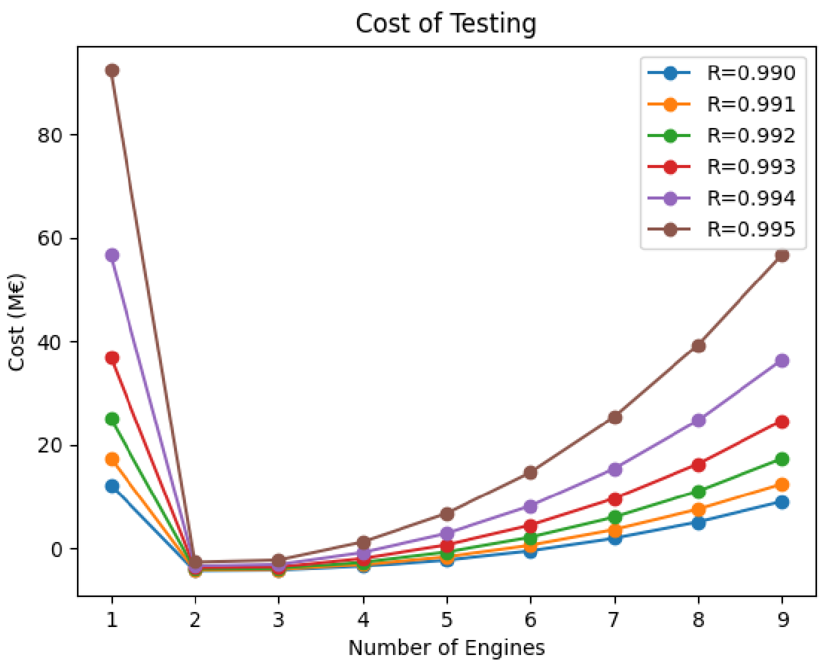

Figure 6.11: Cost of the testing needed to reach reliability targets.

As shown, this cost is proportional to the number of tests. It is worth noting how, despite the different reliability targets being equally spaced, the number of tests, and therefore the cost involved in reaching this target, increases at a higher rate, as the requirement becomes stricter.

This added testing cost includes management and product assurance. Therefore, it impacts other factors in the model, such as the project office costs. Recalculating the development costs, the obtained results are included in figure 6.12 and summarized in table 6.4. Furthermore, the increase in development cost is presented in table 6.5

From the figure, it is possible to see that as the number of engines decreases, initially the cost of development decreases as well, as it is cheaper to develop lighter engines. However, after a certain point, the cost becomes increasing again, as the testing effort increases in order to achieve the reliability targets with a large quantity of engines. There is therefore an optimum in the region between 4 and 6 engines, depending on the reliability target.

The average engine manufacturing costs for 100 flights are calculated from the T1, by considering a learning effect of 90%, and illustrated in figure 6.13. Each flight requires the manufacturing of n engines. This greatly benefits a single engine configuration, as can be seen in the figure. This is due to the fact that, although the engine T1 is lower for an increasing number of engines, the manufacturing

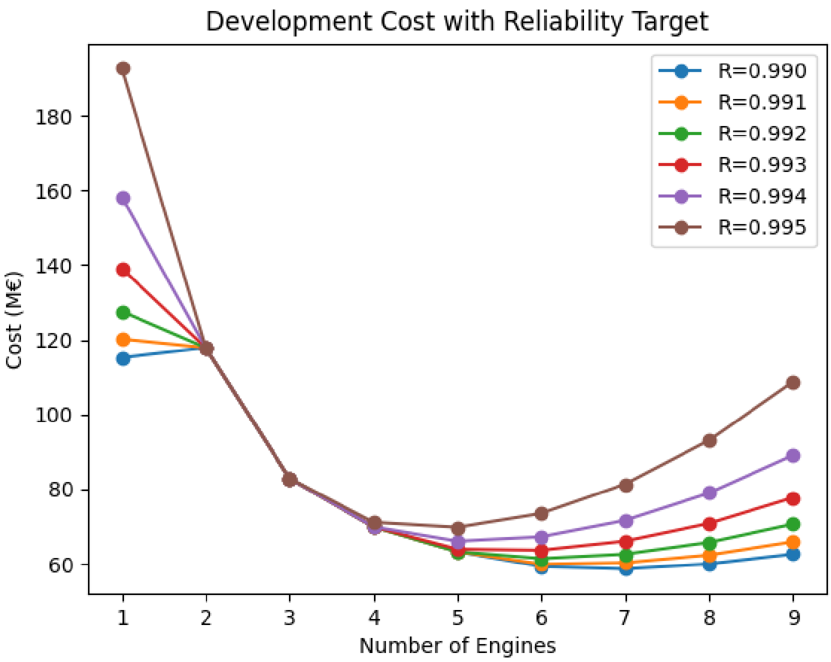

Figure 6.12: Development costs accounting for extra tests needed to reach different reliability targets.

effort of producing a higher quantity of engines dominates the cost result. Furthermore, it is important to point out that the plot is representing the average cost per configuration for 100 flights. The two engine configuration is therefore not as expensive as simply multiplying the single engine configuration costs by two, as the learning curve effect mitigates these costs. Considering only the multi-engine configurations, there is a minimum cost for n = 5.

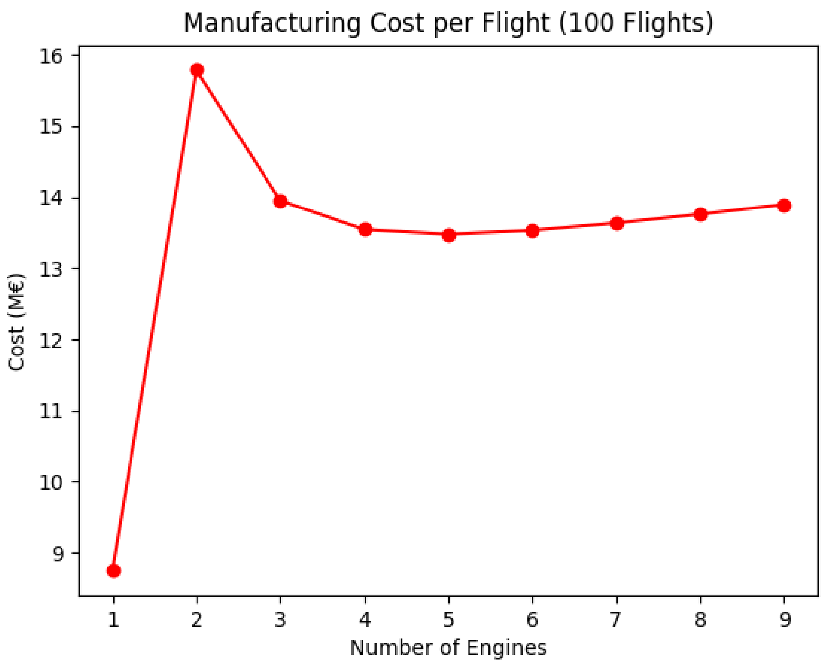

Figure 6.13: Engine manufacturing cost per flight for 100 flights

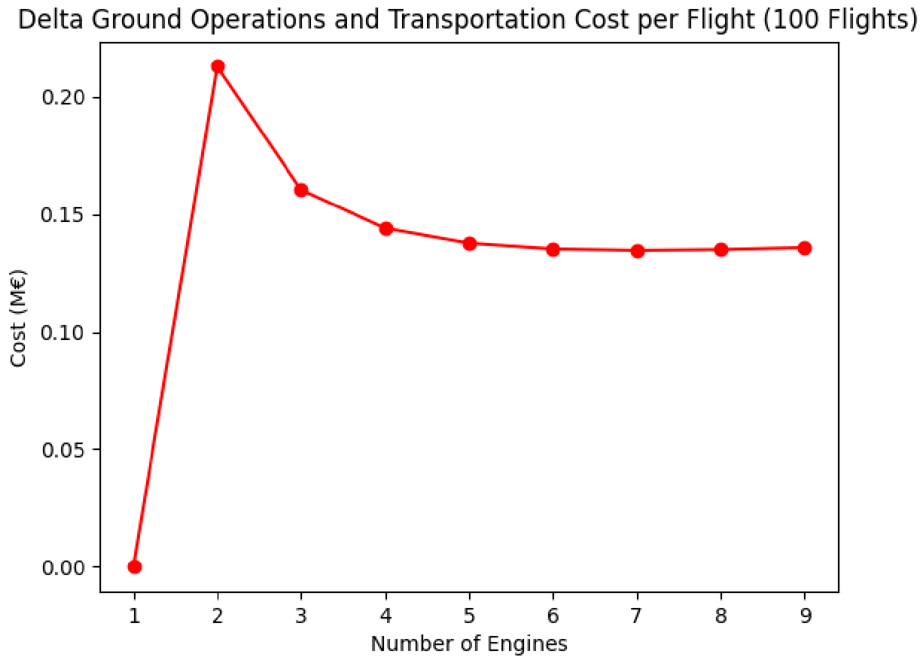

Figure 6.14: Operating cost per flight for 100 flights, pertaining only to the engine cluster.

Regarding operating costs, these will depend on the mass of the vehicle. It is here assumed that there is no extra structural mass needed to support the added engines, that the propellant mass is independent of the configuration, and that there is no increase in mass due to piping, valves, etc. In the future, an integration of this model into an optimization tool will allow to get rid of this assumption and include these mass additions. By simply taking the mass increase coming directly from the cluster mass, it is possible to calculate the operating costs pertaining solely to the engine cluster. The increase in mass influences operating costs in the ground operations (equation 4.13) and transportation domains. This is represented in figure 6.14

The average cost per flight for 100 flights, pertaining solely to the engine subsystem can now be calculated, by adding the amortized engine development costs over the number of flights, the average engine manufacturing costs and the ground and transportation cost fraction dependent on the engines mass. This is represented in figure 6.15 and summarized in table 6.6.

Immediately, it is possible to see in the figure that the single engine configuration is the overall most beneficial configuration in terms of cost. This happens because the lower cost of developing the smaller engines is out-weighted by the increasing cost of producing multiple instances of these engines per flight. In figure 6.16, where costs are broken down according to their origin, it is demonstrated how the steep increase in development cost due to reliability increase translates into a relatively small addition to the total cost per flight, when amortized over the launch vehicles lifetime.

There are some examples that appear to contradict this assertion. For instance, the Saturn V first stage was powered by five F-1 engines. Although there might be other reasons unrelated to cost that motivated this decision, this is also justifiable from a cost and reliability standpoint: Firstly, the F-1 engine is a much more powerful and heavier engine, with a mass of 8,400 kg and a vacuum thrust of 7,770 kN [190]. This make it *a priori* a much more expensive engine to develop, when compared to the Vulcain 2.1 engine. Furthermore, to develop an engine with five times the thrust of this engine, its mass would be much higher, and same for its costs. Considering the relatively low launch rate, as the Saturn V vehicle flew only 13 flights, over 6 years, and purely from an economics standpoint, its amortized development fraction on the cost per flight would be much higher when compared to a multi-engine configuration. Furthermore, it is important to remember that the Saturn V was designed as a manned system, which as it is known implies more stringent reliability requirements. As was shown, a multi-engine configuration sees a reliability increase due to derating and engine-out capability. Other manned systems, such as the Shuttle and the Soyuz, also featured multi-engine configurations.

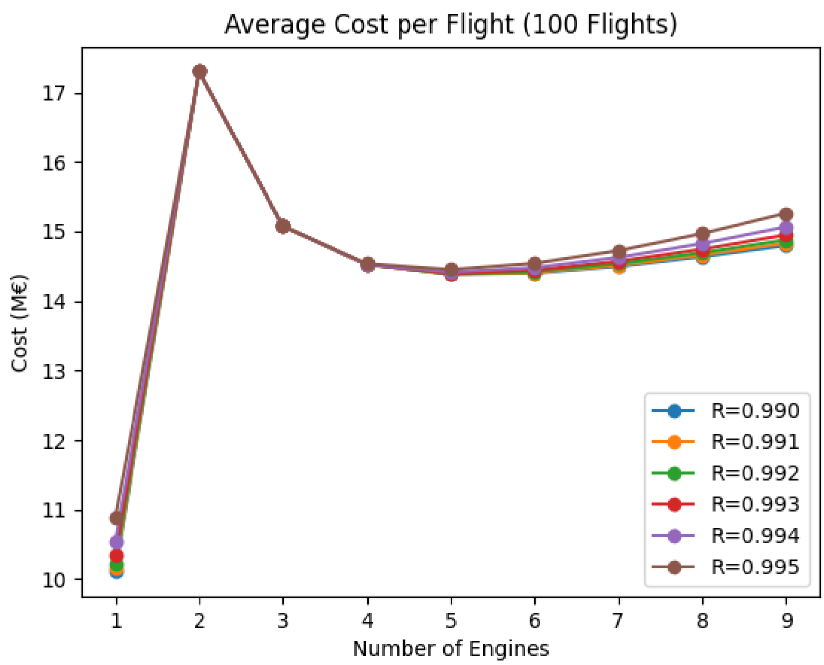

Figure 6.15: Engine manufacturing cost per flight for 100 flights

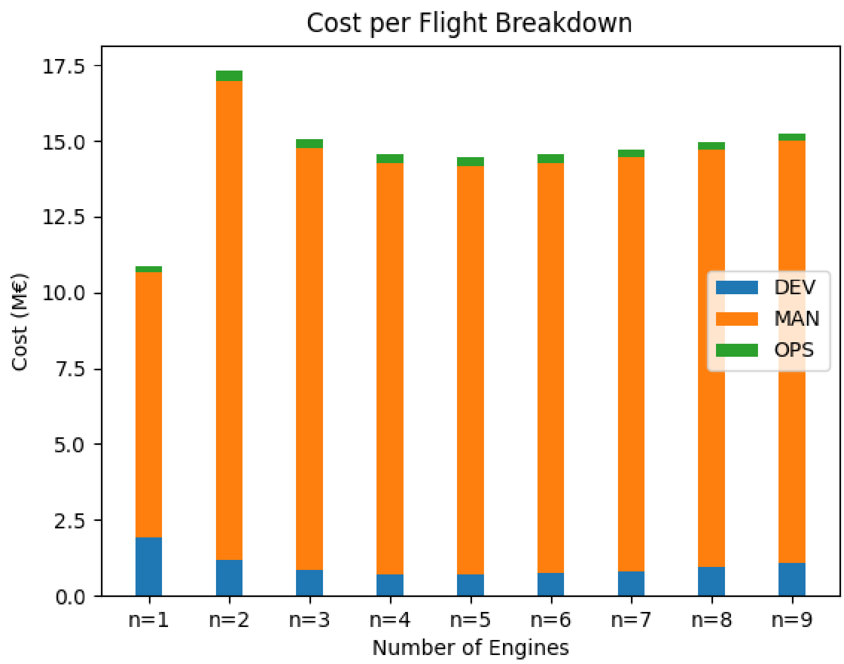

Figure 6.16: Average Cost per Flight breakdown, as a function of the number of engines in the configuration.

Number of Engines	1	2	3	4	5	6	7	8	9
T1	15	15	9	7	6	5	4	4	4
DEV	103	118	83	70	63	59	57	55	54
MAN	9	16	14	14	13	14	14	14	14

OPS	0.21	0.34	0.30	0.29	0.28	0.28	0.28	0.28	0.28

Table 6.3: Summary of costs, in M€ (expendable case).

Number of Engines	1	2	3	4	5	6	7	8	9
DEV' (R=0.990)	115	118	83	70	63	59	59	60	63
DEV' (R=0.991)	120	118	83	70	63	60	60	62	66
DEV' (R=0.992)	128	118	83	70	63	61	63	66	71
DEV' (R=0.993)	139	118	83	70	64	64	66	71	78
DEV' (R=0.994)	158	118	83	70	66	67	72	79	89
DEV' (R=0.995)	193	118	83	71	70	73	81	93	109

Table 6.4: Updated Development Costs, including the cost of testing, in M€.

Number of Engines	1	2	3	4	5	6	7	8	9
DEV/DEV' (R=0.990)	1.11	1.00	1.00	1.00	1.00	1.00	1.03	1.09	1.16
DEV/DEV' (R=0.991)	1.16	1.00	1.00	1.00	1.00	1.01	1.06	1.13	1.22
DEV/DEV' (R=0.992)	1.24	1.00	1.00	1.00	1.00	1.04	1.10	1.19	1.31
DEV/DEV' (R=0.993)	1.35	1.00	1.00	1.00	1.01	1.07	1.16	1.29	1.44
DEV/DEV' (R=0.994)	1.53	1.00	1.00	1.00	1.04	1.13	1.26	1.43	1.65
DEV/DEV' (R=0.995)	1.86	1.00	1.00	1.02	1.10	1.24	1.43	1.69	2.02

Table 6.5: Factor of difference in Development costs.

Number of Engines	1	2	3	4	5	6	7	8	9
CpF (R=0.990)	10.1	17.3	15.1	14.5	14.4	14.4	14.5	14.6	14.8
CpF (R=0.991)	10.2	17.3	15.1	14.5	14.4	14.4	14.5	14.7	14.8
CpF (R=0.992)	10.2	17.3	15.1	14.5	14.4	14.4	14.5	14.7	14.9
CpF (R=0.993)	10.3	17.3	15.1	14.5	14.4	14.4	14.6	14.7	15.0
CpF (R=0.994)	10.5	17.3	15.1	14.5	14.4	14.5	14.6	14.8	15.1
CpF (R=0.995)	10.9	17.3	15.1	14.5	14.5	14.5	14.7	15.0	15.3

Table 6.6: Average Cost per Flight Results, in M€.

6.2.4. Reusability

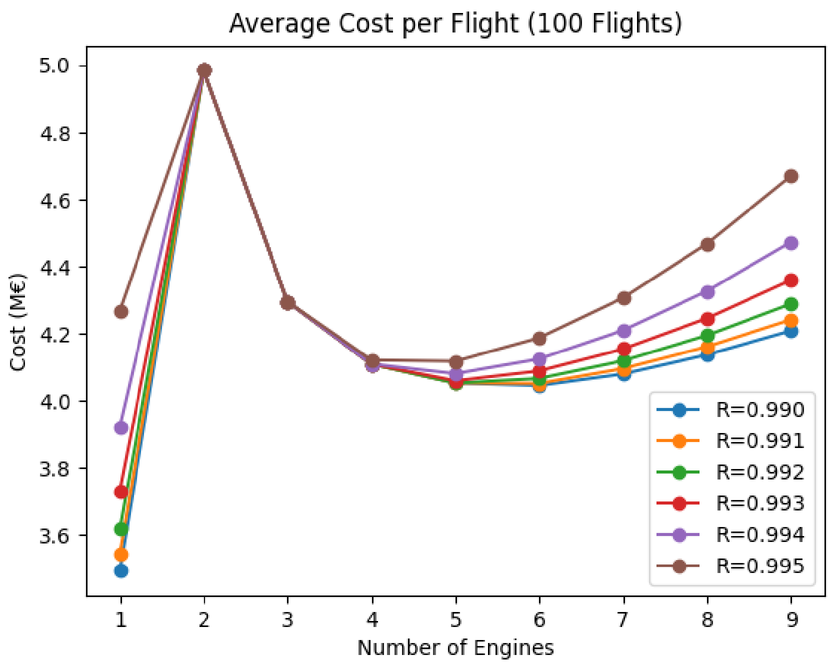

Figure 6.17: Cost per flight including engines and landing gear manufacturing and operations (including retrieval and refurbishment), for different targets of reliability.

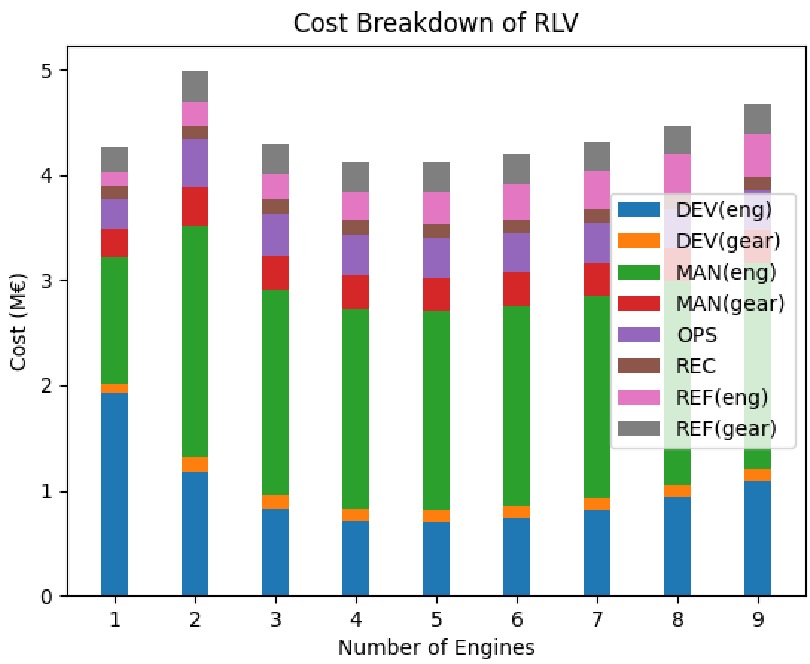

Figure 6.18: Cost breakdown for RLV with engine cluster reliability requirement of 0.995.

Number of Engines	1	2	3	4	5	6	7	8	9
DEV (R=0.990)	115.2	117.8	82.8	69.9	63.2	59.3	58.6	60.0	62.6
DEV (R=0.991)	120.2	117.8	82.8	69.9	63.2	59.9	60.3	62.3	65.9
DEV (R=0.992)	127.6	117.8	82.8	69.9	63.2	61.4	62.6	65.7	70.6
DEV (R=0.993)	139.1	117.8	82.8	69.9	63.9	63.6	66.1	70.8	77.7
DEV (R=0.994)	158.1	117.8	82.8	69.9	66.1	67.2	71.6	78.9	89.1
DEV (R=0.995)	192.6	117.8	82.8	71.1	69.8	73.5	81.3	93.1	108.7
DEV leg	9.4	13.7	12.2	11.8	11.7	11.6	11.6	11.6	11.6
MAN engine (avg)	1.2	2.2	2.0	1.9	1.9	1.9	1.9	1.9	2.0
MAN leg (avg)	0.3	0.4	0.3	0.3	0.3	0.3	0.3	0.3	0.3
Grounds and Transport Costs	0.3	0.5	0.4	0.4	0.4	0.4	0.4	0.4	0.4
Recovery	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Engine Refurbishment	0.1	0.2	0.2	0.3	0.3	0.3	0.4	0.4	0.4
Legs Refurbishment	0.2	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3
CpF (R=0.990)	3.49	4.98	4.30	4.11	4.05	4.05	4.08	4.14	4.21
CpF (R=0.991)	3.54	4.98	4.30	4.11	4.05	4.05	4.10	4.16	4.24
CpF (R=0.992)	3.62	4.98	4.30	4.11	4.05	4.07	4.12	4.20	4.29
CpF (R=0.993)	3.73	4.98	4.30	4.11	4.06	4.09	4.15	4.25	4.36
CpF (R=0.994)	3.92	4.98	4.30	4.11	4.08	4.13	4.21	4.33	4.47
CpF (R=0.995)	4.27	4.98	4.30	4.12	4.12	4.19	4.31	4.47	4.67

Table 6.7: Summary of cost results for the reusability study, in M€.

As the manufacturing costs are dominant in the expendable case, it is expected that in the reusable case these costs are mitigated. Ideally, in order for the reusable solution to be advantageous over the expendable one, the sum of the refurbishment costs of the engines and their retrieval, together with the development, manufacturing and refurbishment of the recovery hardware, would be lower than the average cost of manufacturing the first stage engines every flight of the expendable solution. On top of that, the reusability capability can change the numbers previously obtained, for instance if a propulsive landing is considered, as more propellant needs to be loaded, which leads to a different thrust/mass at engine level, and therefore increased launcher mass. This will be addressed in a sensitivity analysis. As was already mentioned, the cost of manufacturing/refurbishing the rest of the first stage can yield further cost benefits, but this is not being analyzed due to lack of publicly available data. It is also important to mention that there are other reasons, independent from cost, to opt for a multi-engine configuration when considering reusable launch vehicles: particularly when VTVL is employed, having multiple engines with lower thrust is beneficial to land the stage, as using only one engine would require the engine to be throttled to a thrust level that is much lower than its nominal one, in order to softly land just the first stage hardware plus a fraction of remaining propellant.

Firstly, it is assumed that the engines have a lifetime of 10 flights, with a cost aging factor of 115%. After ten flights, the engines aren't further refurbished, and a new cluster is manufactured for the following flight. The same assumptions are applied to the landing system. The results yielded are illustrated in figure 6.17 and summarized in table 6.7.

The first observation is that the single-engine configuration remains the best option for lower reliability levels. However, as the reliability target becomes stricter, the single engine configuration becomes less beneficial, as the amount of testing becomes excessive, and it is more advantageous to switch to a multi-engine configuration, where the derating and engine-out effects contribute to compensate the testing costs. On top of this, the development cost for these smaller engines is also lower, and since the refurbishment costs compensate for the expensive recurring manufacturing costs of the expendable configurations, there are solutions where multiple engine configurations are outperforming the single one in terms of cost. As the number of engines increase and the reliability and cost gains shift from their optimum values, the cost becomes once again increasing with the number of engines. An effect that isn't being considered here is the influence of engine size on the development time, and the indirect effects that that may have on cost. It is expected that the less costly development is accompanied by a reduction in development time, which in turn allows the developer to begin operations sooner, and therefore anticipate the loan payback time, reducing therefore the costs coming from interest rates.

Focusing more on the case where the reliability target is set at R=0.995. Looking at the decomposition of cost into categories, it is possible to note, as expected, how the manufacturing costs lose their weight when switching from an expendable to a reusable configuration.

In figure 6.18, it is possible to see that the added development cost of the landing gear is negligible, and that the retrieval cost corresponds to a small contribution to the overall costs, thanks in part to the learning effect. In this case, the manufacturing costs are not the dominant factors, with a much higher weight being attributed by the refurbishment of both the engines and the landing gear.

6.2.5. Upper Stage Engine Commonality

The possibility of repurposing the first stage engines for use in the upper stage is investigated here. As the Vinci engine has a vacuum thrust of 180 kN, the nine engine configuration, where a thrust of 171 kN was obtained, is a perfect candidate. It is the closest iteration of the multi-engine configurations to the actual Vinci engine, with a 5% error in thrust. This shall be compared to the single-engine configuration, which according to the previous computation is the most advantageous configuration in the expendable case.

Considering an engine mass of 550 kg, cost estimates for development and manufacturing are obtained for the Vinci engine, using the SOLSTICE methodology:

- $T1_{Vinci}$ = 4.85 M€;
- DEV_{Vinci} = 34 M€;

The relevant costs for the Vulcain 2.1 engine (V2.1) and the equivalent engine in a 9 engine configuration ($V2.1_9$) are summarized below:

- $T1_{V2.1}$ = 103 M€;
- $DEV_{V2.1}$ = 15 M€;
- $T1_{V2.1_9}$ = 54 M€;
- $DEV_{V2.1_9}$ = 3.7 M€;

Expendable Launcher

For the single-engine configuration, the average development portion of the cost per flight becomes the sum of the development cost of both engines, amortized over 100 flights:

$$DEV_{n=1} = \frac{DEV_{V2.1} + DEV_{Vinci}}{100} = 1.37 \tag{6.3}$$

The average manufacturing portion is obtained with the T1 production equivalents according to the Crawford learning curve, both until the rank 100.

$$MAN_{n=1_e} = \frac{(T1_{V2.1} + T1_{Vinci}) \cdot \sum_{i=1}^{100} i^{\frac{ln(0.9)}{ln(2)}}}{100} = 11.54 \tag{6.4}$$

$$TOTAL_{no\ commonality_e} = DEV_{n=1} + MAN_{n=1_e} = 12.91 \tag{6.5}$$

On the other hand, for the 9 engine configuration with engine commonality, the development cost is only that of the engine, and the manufacturing cost is simply the learning curve results average, considering 1000 units are produced.

$$DEV_{n=9} = \frac{DEV_{V2.1_9}}{100} = 0.54 \tag{6.6}$$

$$MAN_{n=9_e} = \frac{T1_{V2.1_9} \cdot \sum_{i=1}^{1000} i^{\frac{ln(0.9)}{ln(2)}}}{100} = 15.17 \tag{6.7}$$

$$TOTAL_{commonality_e} = DEV_{n=9} + MAN_{n=9_r} = 15.71 \tag{6.8}$$

With this quick analysis, it is possible to see that although the development cost is lowered by 40%, the manufacturing effort necessary to produce 1000 engines of the same type still far outweighs the effort necessary to produce 100 engines for the first stage and 100 other engines for the upper stage.

Reusable Launcher

If a reusability capability is introduced, according to the principles applied in the previous sections, the development and second stage hardware costs remain the same, however, there are added costs brought by reusability, and the ranks on the learning curve also change.

The added costs related to reusability, REU, consist of the development, manufacturing and refurbishment of landing legs, the refurbishment of the engines and the recovery costs. They are obtained from table 6.7: 0.73 M€ for the single engine configuration, and 0.84 M€ for the nine engine configuration.

For the reusable single-engine configuration with no commonality, the manufacturing costs differ from the expendable case only in that only ten Vulcain 2.1 engines are manufactured to be used in the first stage:

$$MAN_{n=1_r} = \frac{T1_{V2.1} \cdot \sum_{i=1}^{10} i^{\frac{ln(0.9)}{ln(2)}} + T1_{Vinci} \cdot \sum_{i=1}^{100} i^{\frac{ln(0.9)}{ln(2)}}}{100} = 4.02 \tag{6.9}$$

$$TOTAL_{no\ commonality_r} = DEV_{n=1} + MAN_{n=1_r} + REU_{n=1} = 6.12 \tag{6.10}$$

On the other hand, when considering engine commonality of the upper stage with the first stage, 190 Vulcain 2.1_9 engines are manufactured: 90 to be used in the reusable first stage, and 100 to be expended in the upper stage.

$$MAN_{n=9_r} = \frac{T1_{V2.1_9} \cdot \sum_{i=1}^{190} i^{\frac{ln(0.9)}{ln(2)}}}{100} = 3.7 \tag{6.11}$$

$$TOTAL_{commonality_r} = DEV_{n=9} + MAN_{n=9_r} + REU_{n=9} = 5.08 \tag{6.12}$$

This indicates that, if the thrust of the first stage engine allows for it to be used in the upper stage as well, this option leads to a cost reduction, provided that the stage is reusable. A cost reduction is found both in the development and manufacturing costs, however a higher marginal cost due to reusability is expected on the multi-engine configuration, due to the increase in weight. In this example, a cost decrease of 17% was obtained.

The findings of this study are summarized in table 6.8.

Expendable	DEV V2.1	MAN V2.1	DEV Vinci	MAN Vinci	TOTAL
Single-Engine	1.03	8.72	0.34	2.82	12.91
9 engine w/ Commonality	0.54	15.17	-	-	15.71

Reusable	DEV V2.1	MANV2.1	DEV Vinci	MAN Vinci	Reusability	TOTAL
No Commonality	1.03	1.20	0.34	2.82	0.73	6.12
9 engine w/ Commonality	0.54	3.70	-	-	0.84	5.08

Table 6.8: Summary of average cost per flight elements obtained in the commonality study, in M€.

6.2.6. Higher Starting Reliability

In this subsection, the assumed initial reliability is challenged, and it is replaced by the assumption that due to the existing heritage, the apparently insufficient testing campaign is enough to qualify the engine to the same reliability of Vulcain 2, its predecessor.

Assuming a starting reliability of 0.995, the reliability change brought upon by the derating effect and the engine-out capability are actually always beneficial, for the domain of multi-engine configurations considered. This is illustrated in table 6.9.

Number of Engines	1	2	3	4	5	6	7	8	9
Cluster Reliability	0.995	0.9990	0.9989	0.9985	0.9980	0.9975	0.9970	0.9965	0.9959

Table 6.9: Cluster reliability results assuming an original engine reliability of 0.995.

Given these reliability benefits, there is no additional testing required to fulfill the target reliability of 0.995, and therefore no additional testing costs. The single engine remains the most favorable configuration from a cost standpoint as exemplified in table 6.3. The same observation stands for the reusable case.

By adjusting the reliability target, in order to fixate the reliability and have only cost be variable, a similar qualitative result is reached, as in the previous exercise. For the expendable case, the single-engine configuration remains the best choice in the expendable case, whereas in the reusable configuration, the higher the reliability requirement, the more attractive the multi-engine configurations become. This is represented in figure 6.23 and table 6.10.

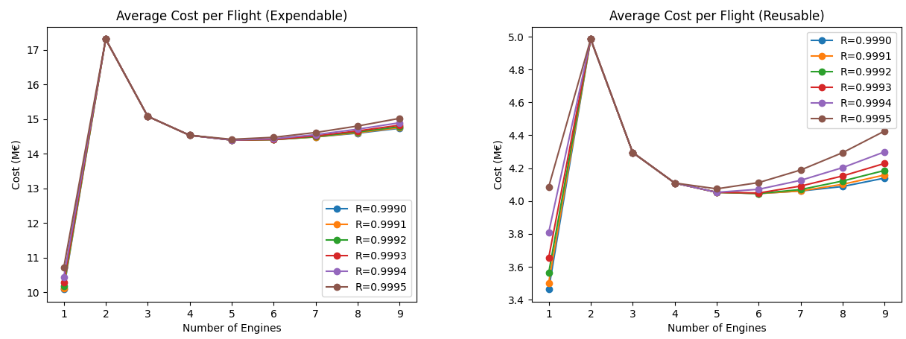

(a) Cost per Flight for initial reliability of 0.995 (expendable case).

(b) Cost per Flight for initial reliability of 0.995 (reusable case).

Figure 6.19: Cost per Flight results considering an initial reliability of 0.995, for the reusable case.

	Number of Engines	1	2	3	4	5	6	7	8	9
Expendable CpF	R=0.9990	10.08	17.30	15.08	14.53	14.39	14.40	14.48	14.59	14.73
	R=0.9991	10.12	17.30	15.08	14.53	14.39	14.40	14.48	14.60	14.75
	R=0.9992	10.18	17.30	15.08	14.53	14.39	14.40	14.49	14.62	14.78
	R=0.9993	10.27	17.30	15.08	14.53	14.39	14.41	14.51	14.65	14.82
	R=0.9994	10.42	17.30	15.08	14.53	14.39	14.43	14.55	14.70	14.89
	R=0.9995	10.70	17.30	15.08	14.53	14.41	14.47	14.61	14.80	15.02
Reusable CpF	R=0.9990	3.46	4.98	4.30	4.11	4.05	4.05	4.06	4.09	4.14
	R=0.9991	3.50	4.98	4.30	4.11	4.05	4.05	4.06	4.10	4.16
	R=0.9992	3.56	4.98	4.30	4.11	4.05	4.05	4.07	4.12	4.19
	R=0.9993	3.66	4.98	4.30	4.11	4.05	4.05	4.09	4.15	4.23
	R=0.9994	3.81	4.98	4.30	4.11	4.05	4.07	4.13	4.20	4.30
	R=0.9995	4.09	4.98	4.30	4.11	4.07	4.11	4.19	4.29	4.43

Table 6.10: Summary of cost results, when considering an initial reliability of 0.995, in M€.

6.2.7. Sensitivity Analysis

Given the big volume of results, for different values of reliability targets and different engine configurations, the sensitivity analysis in this chapter will be done in a qualitative form, by graphically analyzing the effect of varying some key input parameters.

Engine Mass and Structural Mass

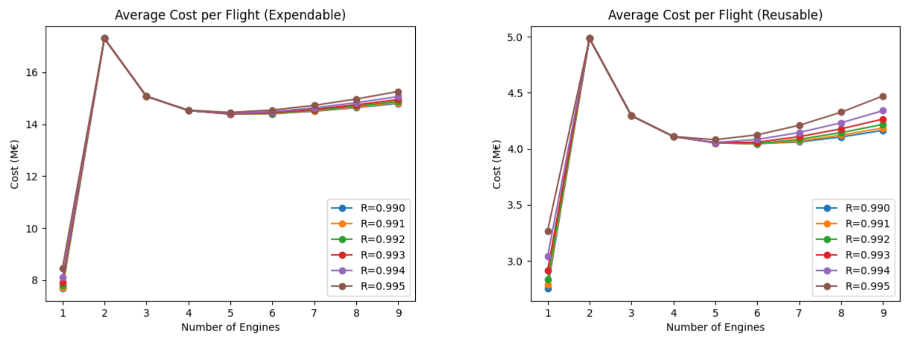

(a) Cost per Flight for 30% lighter original engine (expendable case).

(b) Cost per Flight for 30% lighter original engine (reusable case).

Figure 6.20: Cost per Flight results considering lighter original engine (30%).

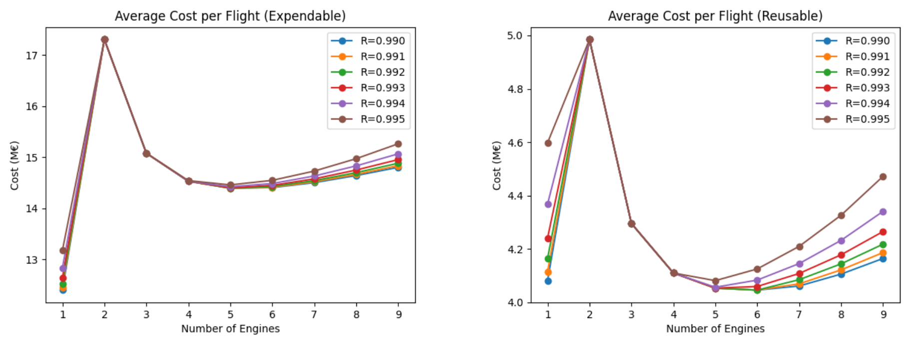

(a) Cost per Flight for 30% heavier original engine (expendable case).

(b) Cost per Flight for 30% heavier original engine (reusable case).

Figure 6.21: Cost per Flight results considering heavier original engine (30%).

Mass is a foundational parameter in this study, and one of the most important variables in anything concerning the space industry, as it has a huge influence on the cost per flight.

The mass of original engine is varied between $\pm$30%. The results are depicted in figure 6.20.

Considering a 30% decrease in mass of the original engine. For the expendable case, it can be seen how the multi-engine configurations become even more distant from the single-engine one, with the best multi-engine options being 75% more expensive than the single-engine one. For the reusable case, it can be seen that the cost per flight curves look very similar in shape to the expendable case, with the single-engine configuration outperforming all the configurations with multiple engines, for all values of cluster reliability.

On the other hand, if an increase of mass of 30% is considered, it can be seen that in the expendable version, multiple engine configurations become more advantageous, however being still beat by the single engine one. For the reusable version, it can be seen that the multi-engine configuration become viable for all reliability targets, with the higher values of reliability displaying a clear benefit when compared to the single-engine configuration.

What can be concluded from this analysis is that as the mass of the starting engine considered for the multi-engine increases, the multi-engine configurations become more advantageous. The opposite is also true, where a lighter engine will benefit more single-engine configurations, for both the expendable and reusable case.

As was already discussed, this model does not take into account the impact on propellant and structural mass coming from the decisions taken in this study. These are especially important in the reusable case, where for propulsive landing it was found that an average of 30% of the landing mass is needed in extra propellant mass [165]. For that reason, the sensitivity of this model to added landing mass is here investigated. Considering an increase of 30% in the mass of the landing legs and in the mass input for the delta-operations, the results depicted in figure 6.22 are obtained.

It can be seen that the added mass leads to a poorer performance of the multi-engine options, when compared to the single-engine configuration. This is expected, as it was seen in figure 6.4 that cluster mass increases with the increase in the number of engines. What this sensitivity analysis result indicates is that this model would benefit from being included in a multi-disciplinary optimization tool, with detailed trajectory, performance and propulsion modules, in order to obtain more granular and accurate results

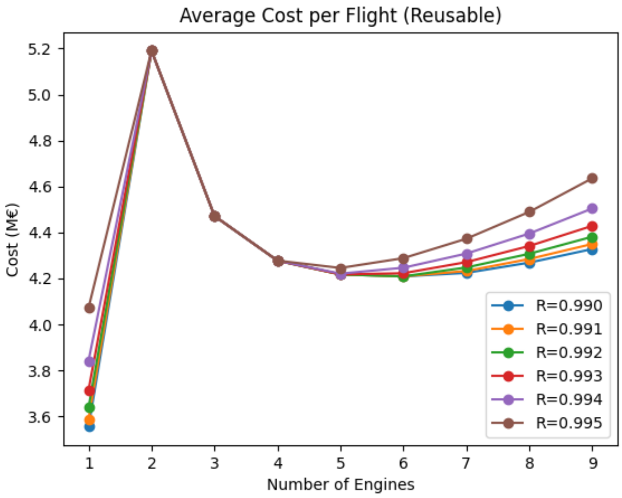

Figure 6.22: Cost per flight sensitivity to 30% higher landing mass.

Catastrophic Failure Percentage

In literature, it was found that catastrophic failure percentage for the engine varies between 10% and 30% [121, 132]. By leaving the best case scenario and opting for a higher catastrophic failure factor of 0.3, as illustrated in figure 6.23, it can be seen that on the expendable case, the resulting effect is that for configurations with higher number of engines, the costs are ”fanned out” and the costs increase for stricter reliability targets, with the general outlook not varying much. On the reusable case, however, this fanning effect is a lot more pronounced, given that the cost curve is itself more flattened, as the cost differences between configurations aren’t that pronounced to start with. This higher catastrophic failure percentage effectively makes configurations with higher numbers of engines unfeasible.

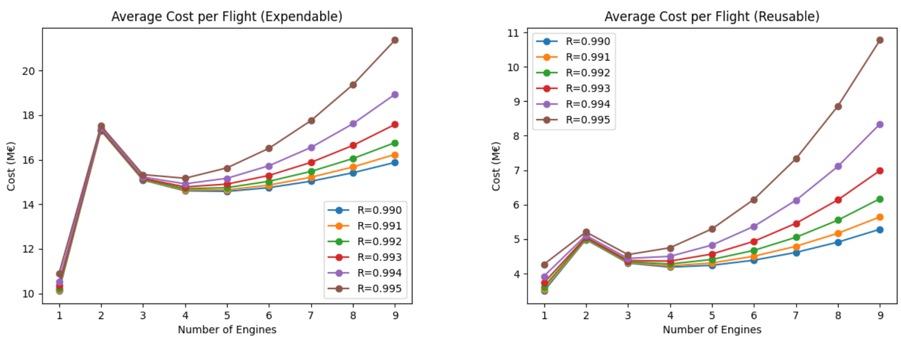

(a) Cost per Flight for CF = 0.3 (expendable case). (b) Cost per Flight for CF = 0.3 (reusable case).

Figure 6.23: Cost per Flight results considering CF = 0.3.

Reliability Growth Factor

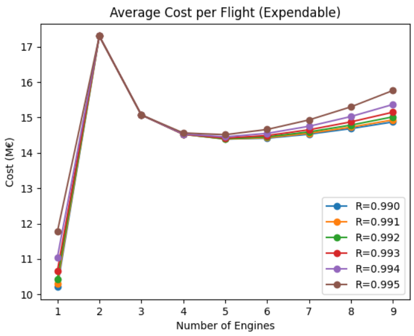

(a) Cost per Flight for $\alpha = 0.36$ (expendable case).

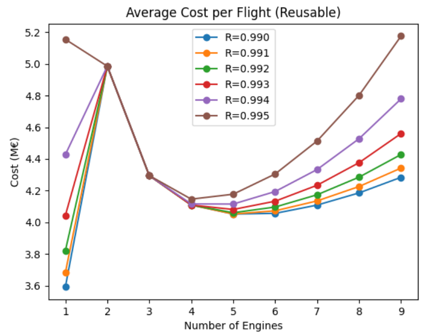

(b) Cost per Flight for $\alpha = 0.36$ (reusable case).

Figure 6.24: Cost per Flight results for $\alpha = 0.36$

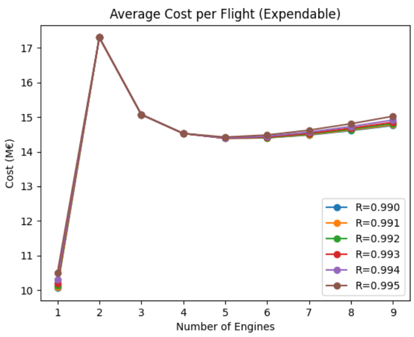

(a) Cost per Flight for $\alpha = 0.44$ (expendable case).

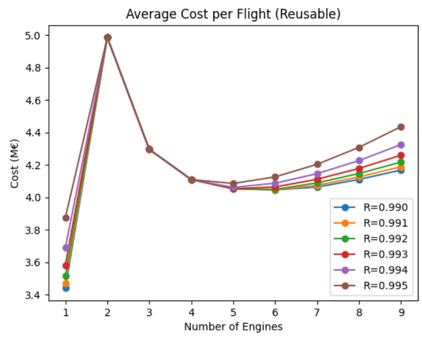

(b) Cost per Flight for $\alpha = 0.44$ (reusable case).

Figure 6.25: Cost per Flight results for $\alpha = 0.44$

The reliability factor dictates the number of tests needed to reach the reliability target. Considering a variation of 10% around the reliability growth factor considered: For the expendable case, once again the result is mostly a fanning of the right tail of the cost curve, not changing the conclusions a lot. For the reusable case however, a lower reliability growth factor makes it so that the costs of testing for the single-engine configuration become a lot more burdensome, which actual makes the multi-engine configurations preferable. On the other hand, a better reliability growth rate benefits the single-engine configuration. These results are depicted in figures 6.26 and 6.27.

Testing Schedule

Using analogous test programs, it was considered an average of 1.25 tests per week (tpw), and with an average of 2.4 ignitions per test, which assuming a four week month, is equivalent to 12 tests per month (tpm). By varying this average amount of tests by half, to 6 and 18 tests per month, the results obtained react similar to the variation in the reliability growth factor. a less intense test campaign harms the single-engine configuration, in comparison to the multi-engine configuration, whereas a cheaper test program resulting from a higher test rate benefits the single-engine configuration. At the same time, the effects are felt much strongly on the reusable case.

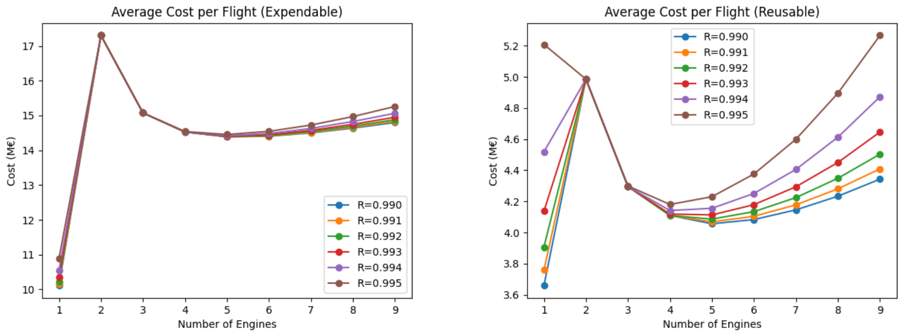

(a) Cost per Flight for tpm = 6 (expendable case). (b) Cost per Flight for tpm = 6 (reusable case).

Figure 6.26: Cost per Flight results for tpm = 6

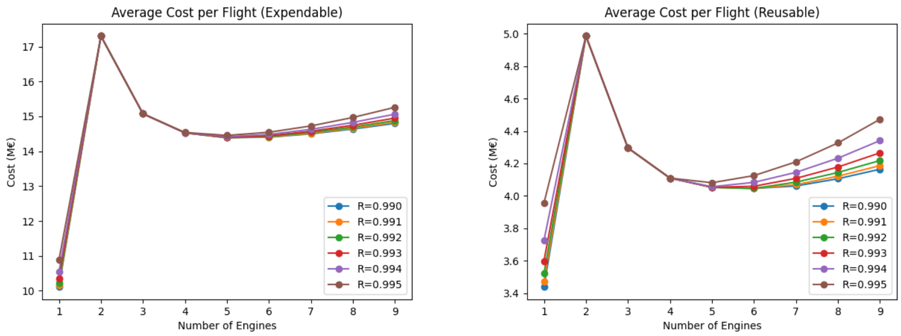

(a) Cost per Flight for tpw = 18 (expendable case). (b) Cost per Flight for tpw = 18 (reusable case).

Figure 6.27: Cost per Flight results for tpw = 18

6.3. Conclusions

The first observation for the expendable case is that at first glance, the single-engine configuration appears to be the most cost effective choice. Despite having a higher development cost, this is effectively amortized over the life time of the vehicle. Furthermore, the advantage of only building or refurbishing one heavier and more complex engine per flight appears to be more beneficial than doing the same for multiple smaller and less complex engines.

A case can be made however for the multi-engine configurations, as in reality the expenditure of money in a project is not homogeneous over time. The high development costs of a bigger engine are a significant barrier to entry to smaller companies with limited funding. The lower development cost sported by the multi-engine configurations will help small commercial companies and start-ups penetrate the market more easily, with the drawback that over the medium and long term, the pricing on their launcher will be a handicap when compared to the more simple and cheaper single-engine configurations.

On top of that, the multi-engine configuration benefits from a natural increase in reliability, as the engine-out capability requirement leads not only to an increase in reliability due to redundancy, but the derated operation also implies a reliability benefit.

Furthermore, when considering the possibility of reusing the first stage, it was found that the feasibility of this solution is highly dependent on the expected refurbishment costs and reliability targets. Assuming a lean refurbishment similar to what SpaceX reportedly does, the multi-engine configurations become more viable as the reliability target for the cluster increases. The optimum number of engines in a given high-reliability configuration appears to be in the 4-5 range. As the number of engines increases beyond that, it becomes less viable to choose one of these configurations, when compared with a single-engine configuration.

Finally, the case where the first stage engine is repurposed for the upper stage as well was studied, and showed that there is a potential for cost savings, especially on the reusable configurations.

It can therefore be concluded that the Ariane 6 design choices make sense, according to this study. As was shown, the recurring costs of manufacturing engines every flight become very burdensome as the number of engines increase, and it is more advantageous to focus on a single engine and qualify it to a high reliability level.

The trend seen in the design of the future reusable launch vehicles is also justified. The cost savings brought about by reusing the first stage and not having to manufacture a high amount of engines every flight are amplified by the cheaper development costs of smaller engines, and the extra non-recurring costs appear to not become a problem cost-wise. Although the optimum number of engines is located around n=5, there is the opportunity to refine this estimate by including the costs of the full launcher, not just the engines, and also by selecting more appropriate reliability and cost factors, informed by profound knowledge of the system, insight into the development, manufacturing and qualification processes, and testing data. All of this can shift the optimum solution and possibly show the benefit of the successful Falcon 9 nine engine configuration. On this note, this study does not find a justification for a 30+ engine configuration. Following the trend lines in cost shown, it would appear that such configuration would result in very high manufacturing costs in the expendable case, compounded further by an expensive qualification process. There is also space to further investigate this subject, as only a 1-engine out capability is contemplated in this study. Having the ability to accommodate for more engine failures would increase the system reliability, and thus bring down the cost of testing. Once there is more information about the Super Heavy system, it will be possible to run a similar study to this one and compare it to the Soviet rocket N1, in order to determine if the system is viable from both a cost and reliability point of view.

7
Conclusion and Recommendations

In this chapter, the research questions presented in the introduction are revisited. It shall be assessed if the answers provided with this research fully address these questions, and how well they do it.

Furthermore, the general findings obtained in this research work are summarized.

Finally, recommendations for further work are provided, in order to advise future research on the topics treated and the avenues this research opened.

7.1. Conclusions

The main question that guided this research project was formulated as such:

Is it possible to estimate the life-cycle costs and expected value of a reusable launch vehicle, taking into account the expected cost of failure?

In order to answer this question, failure had to be incorporated as a cost figure.

Answering first sub-question ”How to estimate the cost of reusable launch vehicle hardware?”, started from a solid foundation in SOLSTICE, which had established a T1 equivalents method to estimate the cost of expendable launch vehicles. This model was complemented with CERs that allowed to estimate the T1 and development costs of recovery hardware. Furthermore, new adjustment factors were developed, based on ESA data, in order to calibrate these CERs to the cases at hand, with an analogy method.

The second sub-question, ”How to estimate the cost of operations of a reusable launch vehicle?”, had as a basis the TRANSCOST model, which was complemented with CERs referring to recovery methods such as in-air capturing, downrange landing as sea, and return to launch site. Furthermore, a refurbishment model was develped, based on CERs with aging and learning effects.

Answering the third sub-question, ”How to model the reliability of a reusable launch vehicle throughout its life-time?”, lead to developing a reliability model, which parametrically estimates the reliability of some components, and relies on historical data for the reliability of other subsystems. The reliability at system level is extrapolated using reliability analysis techniques, specifically fault tree analysis. Combined with reliability growth model analysis, this model allowed to project the reliability of the launch vehicle over its life-cycle. Although there is a high degree of uncertainty, this model lays out a framework for approaching the problem that is as good as the data available to use as input. This means that in the future, as there is more operational data and potentially more data about the testing and qualification program of Falcon 9 and the Merlin engine, the estimates will have a higher confidence.

The process of answering the fourth sub-question, ”How to implement and estimate the cost of reliability increase strategies?”, saw the development of a new CER, directly correlating the time of testing to its cost, based on ESA data points. In conjunction with the reliability model, this allows designers to assess the benefits of an extended test program from both a reliability and a cost standpoint. On top of that, a scenario where the increased reliability due to derating being related to cost was analyzed, in a study about the possible advantages of multi-engine configurations versus single-engine configuration.

The fifth sub-question, ”How does failure factor into the life-cycle cost of a reusable launch vehicle?” lead to the methodology to obtain the expected cost of failure, which for every flight in a launch vehicle life-time provides the expected cost of failure, given the cost of the flight, the probability of failure, and the failure cost.

Finally, the sixth sub-question ”How does the expected value of a reusable launch vehicle vary with reuse?” was answered with the development of a first of its kind value model, which allows designers to estimate the cost of failure at any point during the life-cycle of the launch vehicle, given its cost data, insurance information and investigation procedures.

The expected value model was applied to the Falcon 9 vehicle, and it was discovered that when a reusable launch vehicle is first rolled out, while the learning effects haven't taken off and the system hasn't been given time to mature from a reliability perspective, there is a difficulty in reusing a booster multiple times, as flights beyond the first ones have a negative expected value. This is congruent with expectations taken from the operational experience of the Falcon 9 vehicle, where for the first years of operations, boosters were only flown one to two times before being expended or retired. Furthermore, it was found that given the reliability and cost estimates obtained, the expected value coming from a Falcon 9 flight is actually around a third of the expected profit projected not considering failure. This can be increased to close to 40%, by investing in insurance and reflight-guarantee agreements with the customer. If optimal decisions regarding these insurances are taken according to this model's recommendations, the value can be increased by 3%. Furthermore, it was shown that if the flights with negative value are subsidized, and only the ones with positive value are commercialized, then the expected value is closer to 84% of the expected value. Finally, this figure was reflected on the profit margin, showing that in order to have the expected value equal the expected profit not considering failure, the average profit margin would have to be increased to from 48% to around 75%.

On the other hand, the model developed was applied to a study which sought to assess the potential cost benefits from a single-engine configuration versus a multi-engine configuration, using the Ariane 6 and its Vulcain 2.1 engine as a case-study. It was found that, for expendable launch vehicles, the manufacturing costs dominate the overall cost per flight, and so the development cost savings and reliability benefits coming from opting for a reusable configuration are outweighed by the increasing manufacturing effort from having to produce multiple engines per flight. On the other hand, if manufacturing costs are minimized through for instance the reuse of the first stage, it was found that as the reliability requirement for the propulsion subsystem is increased, the multi-engine configurations become more favorable, eventually being comparable or even out-performing the single-engine configuration from a cost standpoint, for the same reliability targets. When considering this reusable case, it was further discovered that as the number of engines in the first stage and their thrust approaches that of the upper stage engine thrust, considering a commonality strategy leads to cost benefits. In the sensitivity study performed, it was found that for lighter engines, a substitution of the single-engine configuration for a multi-engine one is less advantageous than for heavier engines. Although the expendable configuration still favors the single-engine configuration, as the engine gets heavier, the reusable case holds the multi-engine solutions as more advantageous.

Finally, it is considered that the research project answers the research question postulated at the start of this study, as it was possible to express the probability of failure as a cost figure, be it the expected cost of failure, or the profit margin increase which would lead to the expected value to be equal to the expected profit, not considering failure. However, there is still a need to know how well this model addresses the research question. Due to lack of financial data related to how much money is lost due to a failure, it is difficult to validate this model. This means that in this thesis, a solution has been proposed, but only in the future, if there is more information available, will it be possible to determine if the criterion that assesses how well the research question has been answered is checked out.

7.2. Recommendations

This section will provide advice and recommendations for future researchers, on potential follow up for this work.

First and foremost, the logical next step is to integrate the cost and reliability models developed on an optimization tool, such as the ones developed at TU Delft, like the First Stage Recovery Tool (FRT) by Rozemeijer [165], or the upcoming Masters graduation work by Swati Iyer. Similarly to how Martino's risk and cost models [149] were integrated in the optimization tool developed by Castellini [83], this would allow the algorithm to optimize the launch vehicle configuration based on both reliability and cost, however applied to reusable launch vehicles this time. This is a need that becomes ever more increasing, with multiple VTVL launcher solutions permeating the market and novel reusability concepts, such as in-air capturing being studied and adopted by up and coming companies.

More specifically from a cost perspective, it would be beneficial to develop transparent CERs for the recovery hardware of reusable launch vehicles. As it stands, there are few examples of this type of hardware, and even less information about their mass and cost data. An option would be to gather available information from comparable systems, develop a CER, and use an adjustment factor to calibrate it, according to assumptions coming from cost data and expert judgment.

On top of that, regarding operations cost estimating, the best model currently available for use is TRANSCOST. However, it has documented faults that have been identified in literature, and also has a set of assumptions regarding reusable launch vehicles, which were derived decades ago, and were shown in this work to not correspond to the current scenario of operations for the only reusable launch vehicle currently flying, the Falcon 9. It is suggested that a study on the operation cost estimating for launch vehicles be done, potentially in collaboration with European Space Operations Centre (ESOC) and/or the Guiana Space Centre. Similarly to SOLSTICE, a parametric, transparent, European-based, parametric cost model for operations would allow designers to better estimate the costs coming from launch vehicle operations, based on low-detail information of a system in early phase development.

From a reliability perspective, it will be necessary to in the future tap into more and better data, especially that concerned with the reliability of the landing of VTVL systems. If the designer or launch service provider has testing data, then the model will have a much lower uncertainty. Some avenues of research that this work opens up are the application of more sophisticated methods, such as stochastic processes for parametrization of lifetime distributions. Bayesian methods are recommended, in case it is possible to obtain an informed prior from data.

Still on failure, it is also desirable to study the influence of both infant mortality failures and aging failures on reusable launch systems and their impact in the life-cycle reliability and cost estimates. Furthermore, more sophisticated studies from a reliability, availability, maintainability and safety perspective would enrich the modeling and optimization of launch vehicles. This study has shown a link between reliability and cost. Given the focus of the current commercial launch providers on reliability, side by side with reusability and cost, it seems logic that future optimization work should include a reliability module, as it could uncover insights that would lead to solutions for decreasing costs and optimizing performance.

Concerning the study performed on the advantages of multi-engine configurations, it seems like concepts similar to the N1 rocket or the future Superheavy booster will suffer from high reliability and cost penalties. There are a number of reliability techniques that could be adopted to counter these problems, starting with extending the engine-out capability to tolerate more than one benign failure. This opens up a technical challenge, as well as a cost challenge. This is because in order to justify the recurrent cost of building so many engines per booster, these costs would have to be amortized over multiple flights, which means prolonging the life-time of this hardware. There are therefore solutions needed on the performance of these engines, their reliability and resistance to aging, and their refurbishment. Given that Elon Musk has stated that the Starship launcher could see its cost per flight be reduced to $1M [67], then a study revisiting the boosters with a disproportionate number of engines, compared to the current launchers in the market, as begun by the soviets in the 60s, seems to be once again justified.

Finally, this work also opens the door to future studies in areas other than engineering. From an insurance perspective, it was found in this research that insurance for launch vehicles is at an all time low, and that failures might see an increase in these rates as a response to the market speculation. According to the projected growth in the space sector, translated by a higher launch rate of satellites and a growing occupation of LEO, insurance and risk modeling is a discipline that becomes increasingly more important, as space becomes increasingly more commercial. A more detailed insurance model could be replace the current insurance assumptions taken in the present failure cost model. On top of that, in the field of economics, by having a expected value model, the field of utility theory applied to the current launch vehicle companies could bring new and updated conclusions compared to the ones found in this study, by looking at the commercial companies as entities which are not risk neutral.

Bibliography

[1] Reliable as a kalashnikov assault rifle: the amur methane rocket. https://www.roscosmos.ru/29357/. Accessed: 06-04-2021.

[2] Blue origin makes historic reusable rocket landing in epic test flight. https://www.space.com/31202-blue-origin-historic-private-rocket-landing.html. Accessed: 06-04-2021.

[3] Rocketlab electron. https://www.rocketlabusa.com/rockets/electron/. Accessed: 06-04-2021.

[4] Spacex falcon 9. https://www.spacex.com/vehicles/falcon-9/. Accessed: 22-06-2021.

[5] Elon musk touts low cost to insure spacex rockets as edge over competitors. https://www.cnbc.com/2020/04/16/elon-musk-spacex-falcon-9-rocket-over-a-million-dollars-less-to-insure.html, . Accessed: 15/02/2022.

[6] The hangar / falcon 9 - specifications. https://www.spaceflightinsider.com/hangar/falcon-9/, . Accessed: 18/04/2022.

[7] Secretary, dos / chairman, isro's new year message for 2021 "laying a firm foundation for the next decade". https://www.isro.gov.in/update/01-jan-2021/secretary-dos-chairman-isro%E2%80%99s-new-year-message-2021-%E2%80%9Claying-firm-foundation-next. Accessed: 06-04-2021.

[8] Spacex launch monday will signify a new advance in reusing rockets. https://spaceflightnow.com/2018/12/01/spacex-launch-sunday-will-signify-a-new-advance-in-reus Accessed: 06/04/2022.

[9] China to test rocket reusability with planned long march 8 launcher. https://spacenews.com/china-to-test-rocket-reusability-with-planned-long-march-8-launcher/. Accessed: 06-04-2021.

[10] Space rider: Europe's first reusable space transportation system. https://www.esa.int/Enabling_Support/Space_Transportation/Space_Rider, . Accessed: 06-04-2021.

[11] Spacex. https://www.spacex.com/, . Accessed: 06-04-2021.

[12] Esa plans demonstration of a reusable rocket stage. https://www.esa.int/Enabling_Support/Space_Transportation/ESA_plans_demonstration_of_a_reusable_rocket_stage. Accessed: 06-04-2021.

[13] 4cost aces. https://www.4cost.de/en/software-2/aces-2/. Accessed: 23/06/2021.

[14] Airbus unveils 'adeline' re-usable rocket concept. https://www.bbc.com/news/science-environment-33006056. Accessed: 15/04/2021.

[15] Arianespace flight 157 - inquiry board submits findings. https://www.esa.int/Enabling_Support/Space_Transportation/Arianespace_Flight_157_-_Inquiry_Board_submits_findings2, . Accessed: 12/05/2022.

[16] Investigators say erroneous navigation input led ariane 5 rocket off course. https://spaceflightnow.com/2018/02/23/investigators-say-erroneous-navigation-input-led-ariane-5-rocket-off-course/, . Accessed: 12/05/2022.

[17] N° 33–1996: Ariane 501 - presentation of inquiry board report. https://www.esa.int/Newsroom/Press_Releases/Ariane_501_-_Presentation_of_Inquiry_Board_report, . Accessed: 12/05/2022.

[18] Ariane 5 failure investigation focuses on upper stage. https://www.spaceflightnow.com/ariane/v142/010713followup.html, . Accessed: 12/05/2022.

[19] Ariane 6 on the way to flight. https://www.esa.int/Enabling_Support/Space_Transportation/Ariane/Ariane_6_on_the_way_to_flight, . Accessed: 14/03/2022.

[20] List of ariane launches. https://en.wikipedia.org/wiki/List_of_Ariane_launches, . Accessed: 12/05/2022.

[21] Arianespace flight 142: Inquiry board report and action plan. https://sma.nasa.gov/LaunchVehicle/assets/arianespace-flight-142-inuiry-board-report.pdf, . Accessed: 12/05/2022.

[22] Air force budget reveals how much spacex undercuts launch prices. https://arstechnica.com/science/2017/06/air-force-budget-reveals-how-much-spacex-undercuts-launch-prices/. Accessed: 02/05/2022.

[23] Podcast: Interview with spacex's elon musk. ://aviationweek.com/defense-space/space/podcast-interview-spacexs-elon-musk. Accessed: 06/04/2022.

[24] Full thrust on europe's new ariane 6 rocket. https://www.bbc.com/news/science-environment-40366736. Accessed: 12/05/2022.

[25] The new rockets racing to make space affordable. https://www.bloomberg.com/graphics/2018-rocket-cost/. Accessed: 02/05/2022.

[26] Elon musk's spacex wins 2 pentagon contracts for nearly $160 million to launch missions with its falcon 9 rockets. https://www.businessinsider.com/elon-musk-spacex-won-two-pentagon-contracts-launch-military-rockets-2021-3?international=true&r=US&IR=T. Accessed: 02/05/2022.

[27] Spacex dragon. https://www.spacex.com/vehicles/dragon/. Accessed: 14-04-2021.

[28] Here's everything elon musk told reporters about the reusable rocket that will fly twice within 24 hours. https://www.cnbc.com/2018/05/11/full-elon-musk-transcript-about-spacex-falcon-9-block-5.html, . Accessed: 28/04/2022.

[29] Elon musk tweet, 19/08/2020. https://www.teslarati.com/spacex-15m-florida-falcon-refurbishment/, . Accessed: 06/04/2022.

[30] Elon musk tweet, 19/08/2020. https://twitter.com/elonmusk/status/1295883862380294144?s=20&t=mPHDOxV97zroispkzFE5dg, . Accessed: 02/05/2022.

[31] How much does it cost to launch a reused falcon 9? elon musk explains why reusability is worth it. https://www.elonx.net/how-much-does-it-cost-to-launch-a-reused-falcon-9-elon-musk-explains-why-reusab
. Accessed: 02/05/2022.

[32] Ariane 6. https://www.esa.int/Enabling_Support/Space_Transportation/Launch_vehicles/Ariane_6, note = Accessed: 04/04/2022.

[33] Spacex getting better reusability with vastly improved and expanded heat shielding. https://www.nextbigfuture.com/2018/02/spacex-getting-better-reusability-with-vastly-improved-and-expanded-heat-shield
html, . Accessed: 14/03/2022.

[34] Musk details block 5 improvements to falcon 9. https://spacenews.com/musk-details-block-5-improvements-to-falcon-9/, . Accessed: 14/03/2022.

[35] Falcon 9 launch vehicle. https://web.archive.org/web/20181012214449/http://www.spaceflight101.net/falcon-9-launch-vehicle-information.html. Accessed: 18/04/2022.

[36] Why the us can beat chine: The facts about spacex costs. https://web.archive.org/web/20130328121051/http://www.spacex.com/usa.php. Accessed: 18/04/2022.

[37] Spacex plans huge expansion for floridan falcon refurbishment facilities. https://www.teslarati.com/spacex-florida-expansion-falcon-heavy-refurbishment-facility/, . Accessed: 06/04/2022.

[38] Spacex to receive $15m from florida to build falcon refurbishment facility. https://www.teslarati.com/spacex-15m-florida-falcon-refurbishment/, . Accessed: 06/04/2022.

[39] Europe complains: Spacex rocket prices are too cheap to beat. https://www.fool.com/investing/2020/11/10/europe-space-rocket-incredibly-expensive-airbus/. Accessed: 14/03/2022.

[40] When it comes to military launches, spacex may no longer be the low-cost provider. https://www.forbes.com/sites/lorenthompson/2020/09/17/when-it-comes-to-military-launches-spacex-may-no-longer-be-the-low-cost-provider/?sh=1e9d66494c3b. Accessed: 02/05/2022.

[41] Space insurers warn that current low rates are not sustainable. https://spacenews.com/space-insurers-warn-that-current-low-rates-are-not-sustainable/. Accessed: 15-02-2022.

[42] Spacex: Elon musk breaks down the cost of reusable rockets. https://www.inverse.com/innovation/spacex-elon-musk-falcon-9-economics. Accessed: 02/05/2022.

[43] i-space. http://www.i-space.com.cn/. Accessed: 15/04/2021.

[44] Commercial space transportation licenses. https://www.faa.gov/space/licenses. Accessed: 12/05/2022.

[45] List of falcon 9 first-stage boosters. https://en.wikipedia.org/wiki/List_of_Falcon_9_first-stage_boosters, . Accessed: 12/05/2022.

[46] List of falcon 9 and falcon heavy launches. https://en.wikipedia.org/wiki/List_of_Falcon_9_and_Falcon_Heavy_launches, . Accessed: 12/05/2022.

[47] N1 soviet launch vehicle. https://www.britannica.com/technology/N1. Accessed: 12/05/2022.

[48] Neutron. https://www.rocketlabusa.com/rockets/neutron/, . Accessed: 15/04/2021.

[49] Rocket lab reveals neutron launch vehicle's advanced architecture. https://www.rocketlabusa.com/updates/rocket-lab-reveals-neutron-launch-vehicles-advanced-architecture/, . Accessed: 15/04/2022.

[50] New glenn. https://www.blueorigin.com/new-glenn/, . Accessed: 15/04/2021.

[51] New shepard. https://www.blueorigin.com/new-shepard/, . Accessed: 15/04/2021.

[52] Select committee on science and technology minutes of evidence - examination of witness (questions 120-128). https://publications.parliament.uk/pa/cm200304/cmselect/cmsctech/711/4070509.htm. Accessed: 15/02/2022.

[53] Pcec – project cost estimating capability. https://www.nasa.gov/offices/ocfo/functions/models_tools/PCEC. Accessed: 12/05/2022.

[54] Price systems. https://www.pricesystems.com/. Accessed: 23/06/2021.

[55] What is censored data. https://reliability.readthedocs.io/en/latest/What%20is%20censored%20data.html. Accessed: 25/04/2022.

[56] Space insurance costs to rocket after satellite crash. https://www.reuters.com/article/space-insurance/space-insurance-costs-to-rocket-after-satellite-crash-idINKCN1U(
edition-redirect=in. Accessed: 15/02/2022.

[57] Rocket lab successfully deploys 34 satellites and catches rocket booster returning from space with helicopter. https://www.rocketlabusa.com/updates/new-blog-post-8/. Accessed: 12/05/2022.

[58] Rlv-td. https://www.isro.gov.in/launcher/rlv-td. Accessed: 15/04/2021.

[59] Rocket lab progresses flight review, recovers first stage following successful ocean splashdown). https://www.rocketlabusa.com/about-us/updates/rocket-lab-progresses-flight-review-recovers-first-stage-following-successful-o
Accessed: 17-05-2021.

[60] Galorath. https://galorath.com/. Accessed: 23/06/2021.

[61] List of space shuttle missions. https://en.wikipedia.org/wiki/List_of_Space_Shuttle_missions. Accessed: 12/05/2022.

[62] Spacex explains why the u.s. space force is paying $316 million for a single launch. https://spacenews.com/spacex-explains-why-the-u-s-space-force-is-paying-316-million-for-a-single-laun
. Accessed: 02/05/2022.

[63] Learn - virgin galactic. https://www.virgingalactic.com/learn/, . Accessed: 15/04/2021.

[64] Launch anomaly sidelines falcon 9, investigation begins. https://www.satellitetoday.com/business/2015/06/29/launch-anomaly-sidelines-falcon-9-investigation-begins/,
. Accessed: 12/05/2022.

[65] Picking up the pieces – spacex begins investigation into falcon 9 on-pad explosion. https://spaceflight101.com/spacex-falcon-9-failure-investigation-after-amos-6-loss/, . Accessed: 12/05/2022.

[66] Spacex capabilities and services. https://www.spacex.com/media/Capabilities&Services.pdf, . Accessed: 02/05/2022.

[67] Elon musk: Each starship launch could cost just $1 million. https://futurism.com/elon-musk-starship-launch-1-million. Accessed: 12/05/2022.

[68] Spacex starship. https://www.spacex.com/vehicles/starship/. Accessed: 14-04-2021.

[69] The cost of trucking. https://www.imiproducts.com/blog/the-cost-of-trucking/. Accessed: 07-03-2022.

[70] Profile of michel van pelt on researchgate. https://www.researchgate.net/profile/Michel-Van-Pelt-2. Accessed: 12/05/2022.

[71] N° 13–2019: Vega flight vv15: findings of the independent inquiry commission's investigations. https://www.esa.int/Newsroom/Press_Releases/Vega_flight_VV15_findings_of_the_Independent_Inquiry_Commission_s_investigations, . Accessed: 12/05/2022.

[72] N° 23–2020: Vega flight vv17 failure: Arianespace and esa appoint an independent inquiry commission. https://www.esa.int/Newsroom/Press_Releases/Vega_flight_VV17_failure_Arianespace_and_ESA_appoint_an_independent_Inquiry_Commission, . Accessed: 12/05/2022.

[73] NASA / Air Force Cost Model. 2002.

[74] David L. Akin. ENAE 791 - Launch and Entry Vehicle Design Lecture Notes - Cost Estimation and Engineering Economics. 2016.

[75] David L. Akin. ENAE 791 - Launch and Entry Vehicle Design Lecture Notes - Mass Estimating Relationships. 2016.

[76] David L. Akin. ENAE 791 - Launch and Entry Vehicle Design Lecture Notes - Reliability, Redundancy, and Resiliency. 2016.

[77] Timothy P. Anderson. The national estimator - learning curve primer part 1 - unit theory. 2011.

[78] Arianegroup. SOLUTIONS FOR COMMERCIAL AND SCIENTIFIC PAYLOADS. pages 6–7, 2019.

[79] Kenneth J. Arrow and Robert C. Lind. Uncertainty and the Evaluation of Public Investment Decisions. *Journal of Natural Resources Policy Research*, 6(1):29–44, 2014. ISSN 19390459. doi: 10.1080/19390459.2014.867640.

[80] Lucía Ayala, Fernández Carsten, and Wiedemann Vitali. Analysis of Space Launch Vehicle Failures and Post □ Mission Disposal Statistics. *Aerotecnica Missili Spazio*, (0123456789), 2022. ISSN 2524-6968. doi: 10.1007/s42496-022-00118-5. URL https://doi.org/10.1007/s42496-022-00118-5.

[81] Malaya Kumar Biswal M, Ramesh Kumar V, and Noor Basanta Das. A Comparative Study on Orbital Launch Systems for Human Mission to Moon and Mars. pages 0–3, 2021. doi: 10.2514/6.2021-3274.

[82] Jasper Bouwmeester. The architecture of cubesats and pocketqubes towards a lean and reliable implementation of subsystems and their interfaces.

[83] Francesco Castellini. Multidisciplinary Design Optimization For Expendable Launch Vehicles. page 179, 2012. URL https://www.politesi.polimi.it/handle/10589/56841.

[84] Tao Chen, Jiawen Li, Ping Jin, and Guobiao Cai. Reusable rocket engine preventive maintenance scheduling using genetic algorithm. *Reliability Engineering and System Safety*, 114(1):52–60, 2013. ISSN 09518320. doi: 10.1016/j.ress.2012.12.020. URL http://dx.doi.org/10.1016/j.ress.2012.12.020.

[85] Paul D Collopy. Aerospace System Value Models : A Survey and Observations. (September 2009), 2015. doi: 10.2514/6.2009-6560.

[86] Stephane Contant. Design and Optimization of a Small Reusable Launch Vehicle Using Vertical Landing Techniques. page 133, 2019.

[87] Futron Corporation. Design Reliability Comparison for SpaceX Falcon Vehicles. 3202 (November), 2004.

[88] R Curran, S Raghunathan, and M Price. Review of aerospace engineering cost modelling: The genetic causal approach. *Progress in Aerospace Sciences*, 40(8):487–534, 2004. ISSN 0376-0421. doi: http://dx.doi.org/10.1016/j.paerosci.2004.10.001. URL http://www.sciencedirect.com/science/article/pii/S0376042104000594.

[89] elseco Curtis Monier. Expert judgment consultation - personal communication.

[90] D. J. Davis. An Analysis of Some Failure Data. *Journal of the American Statistical Association*, 47(258):113–150, 1952. ISSN 1537274X. doi: 10.1080/01621459.1952.10501160.

[91] J-F. Delange, P. Alliot, A. Lekeux, and B. Vieille. VINCI ® , the European reference for Ariane 6 upper stage cryogenic propulsive system. 2015. doi: 10.2514/6.2015-4063.

[92] Department of Defense of the USA. Military Handbook Reliability Growth Management. (I), 1981.

[93] Department of Defense of the USA. Reliability Prediction of Electronic Equipment. *Military Handbook MIL-HDBK-217F*, page 205, 1991.

[94] Balbir S. Dhillon. Reliability growth: A survey. *Microelectronics Reliability*, 20(5):743–754, 1980. ISSN 00262714. doi: 10.1016/0026-2714(80)90403-5.

[95] Casey Dreier. An Improved Cost Analysis of the Apollo Program. *Space Policy*, (xxxx):101476, 2022. ISSN 02659646. doi: 10.1016/j.spacepol.2022.101476. URL https://doi.org/10.1016/j.spacepol.2022.101476.

[96] N. T. Drenthe, B. T.C. Zandbergen, R. Curran, and M. O. Van Pelt. Cost estimating of commercial smallsat launch vehicles. *Acta Astronautica*, 155:160–169, 2019. ISSN 00945765. doi: 10.1016/j.actaastro.2018.11.054.

[97] Nigel Drenthe. Expert judgment consultation - personal communication.

[98] Nigel Drenthe. Falcon 9 Cost Estimate with SOLSTICE. 2016.

[99] Nigel. T Drenthe. Small Orbital Launch Systems, a Tentative Initial Cost Estimate. 2016. URL http://repository.tudelft.nl/.

[100] J T Duane. LEARNING CURVE APPROACH TO RELIABILITY MONITORING. *IEEE TRANSACTIONS ON AEROSPACE*, 1964. ISSN 15206882. doi: 10.1021/ac60048a710.

[101] Paul Meier E.L. Kaplan. Nonparametric Estimation from Incomplete Observations Author (s): E . L . Kaplan and Paul Meier Source : Journal of the American Statistical Association , Vol . 53 , No . 282 (Jun ., 1958), pp . 457- Published by : American Statistical Association Sta. *American statistical Association*, 53(282):457–481, 1958. ISSN 01621459.

[102] Mohammad Emami and Seyed Hossein Iranmanesh. EVALUATING THE RELIABILITY OF WORK BREAK. (February), 2013. doi: 10.13140/2.1.4164.2569.

[103] KT Engineering. Operations cost model version 1.1.

[104] EPSMA. Guidelines to Understanding Reliability Prediction. (June):1–29, 2005.

[105] Ruwan Ernst. Liquid Rocket Analysis (LiRA).

[106] NASA ESAS. Nasa's exploration systems architecture study.

[107] European Cooperation for Space Standardization - ECSS. Technology readiness level (TRL) guidelines. (March):62, 2017.

[108] Chris Everett and Homayoon Dezfuli. A BAYESIAN APPROACH TO RELIABILITY GROWTH MODELING. (Psam 13):1–12, 2016.

[109] FAA. Special Report : Update of the Space and Launch Insurance Industry. 1998.

[110] FAA. Commercial Space and Launch Insurance : Current Market and Future Outlook Fourth Quarter 2002 Quarterly Launch Report. 2002.

[111] William Farr. Software Reliability Modeling Survey.pdf, 1983.

[112] Federal Aviation Administration. Guide to Reusable Launch and Reentry Vehicle Reliability Analysis. (April), 2005.

[113] Joseph R. Fragola. How safe must a potential crewed launcher be demonstrated to be before it is crewed? *Journal of Loss Prevention in the Process Industries*, 22(5):657–663, 2009. ISSN 09504230. doi: 10.1016/j.jlp.2009.05.009. URL http://dx.doi.org/10.1016/j.jlp.2009.05.009.

[114] Dan M. Frangopol. Life-cycle cost design of deteriorating structures.

[115] Dan M Frangopol, Maarten-Jan Kallen, and Jan M van Noortwijk. Probabilistic models for life□cycle performance of deteriorating: Review and Future Directions, 2004.

[116] A. Fries and A. Sen. Survey of Discrete Reliability-Growth Models. 45(4):582–604, 1996.

[117] Susie Go. A historical survey with success and maturity estimates of launch systems with RL10 upper stage engines. *Proceedings - Annual Reliability and Maintainability Symposium*, pages 491–495, 2008. ISSN 0149144X. doi: 10.1109/RAMS.2008.4925845.

[118] Major Greenwood. A report on the natural duration of cancer.

[119] Geoffrey Grimmett and Dominic Welsh. Probability: An Introduction. 2014.

[120] Todd Herrmann. A Critical Parameter Optimization of Launch Vehicle Costs. 2006.

[121] Zhaofeng Huang, Jeffry A. Fint, and Frederick M. Kuck. Key reliability drivers of liquid propulsion engines and a reliability model for sensitivity analysis. *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, (July):1–12, 2005. doi: 10.2514/6.2005-4436.

[122] Edward L. Keith. Advanced reusable launch vehicle science, analysis and design. 2004.

[123] Shari Keller. DEVELOPMENT APPLICATION OF A SPACE LAUNCH SYSTEM VALUE MODEL TO A LUNAR HELIUM-3 MINING MISSION. 2013.

[124] Shari Keller and Paul Collopy. Value modeling for a space launch system. *Procedia Computer Science*, 16:1152–1160, 2013. ISSN 18770509. doi: 10.1016/j.procs.2013.01.121. URL http://dx.doi.org/10.1016/j.procs.2013.01.121.

[125] Kyungmee O. Kim, Taeseong Roh, Jae Woo Lee, and Ming J. Zuo. Derating design for optimizing reliability and cost with an application to liquid rocket engines. *Reliability Engineering and System Safety*, 146:13–20, 2016. ISSN 09518320. doi: 10.1016/j.ress.2015.10.005. URL http://dx.doi.org/10.1016/j.ress.2015.10.005.

[126] D. Koelle. Handbook of cost engineering for space transportation systems with transcost 7.2 : statistical-analytical model for cost estimation and economical optimization of launch vehicles. 2007.

[127] D. E. Koelle. Cost-optimal size of future reusable launch vehicles. *Acta Astronautica*, 47(2): 205–213, 2000. ISSN 00945765. doi: 10.1016/S0094-5765(00)00060-6.

[128] Dietrich E. Koelle. Cost efficiency as design and selection criterion for future launch vehicles. *Acta Astronautica*, 57(2-8):623–629, 2005. ISSN 00945765. doi: 10.1016/j.actaastro.2005.02.001.

[129] Dietrich E Koelle. *Handbook of Cost Engineering and Design of Space Transportation Systems*. Number TCS-TR-200. TransCostSystems (TCS), Ottobrunn Germany, 8.2 edition, 2013.

[130] Jung S. Kong. Cost–reliability interaction in life-cycle cost optimization of deteriorating structures.

[131] Zuzana Krajcusková. A brief survey of reliability growth models. *Proc. of 13th International Czech-Slovak Scientific Conference “RADIOELEKTRONIKA*, 2003.

[132] Zachary C Krevor. A Methodology to Link Cost and Reliability for Launch Vehicle Design A Methodology to Link Cost and Reliability for Launch Vehicle Design. *Aerospace Engineering*, (August):1–264, 2007. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.472.2487&rep=rep1&type=pdf.

[133] Bernard F. Kutter. Commercial launch services: An enabler for launch vehicle evolution and cost reduction. *Collection of Technical Papers - Space 2006 Conference*, 1(September):664–672, 2006. doi: 10.2514/6.2006-7271.

[134] Juseong Lee and Mihaela Mitici. Multi-objective Analysis of Condition-based Aircraft Maintenance Strategies using Discrete Event Simulation. *Proceedings - Annual Reliability and Maintainability Symposium*, 2021-May:1–6, 2021. ISSN 0149144X. doi: 10.1109/RAMS48097.2021.9605761.

[135] Lawrence M. Leemis. *Reliability - Probabilistic Models and Statistical Methods*. Prentice Hall, Englewood Cliffs, N.J., second edition, 2009. ISBN 978-0-692-00027-4.

[136] Charles Lillie and Bruce Thompson. Parametric cost estimation for space science missions - art. no. 701827. 07 2008. doi: 10.1117/12.789615.

[137] Eric Mahr, Anh Tu, and Anil Gupta. Development of the Small Satellite Cost Model 2019 (SSCM19). *IEEE Aerospace Conference Proceedings*, 2019, 2020. ISSN 1095323X. doi: 10.1109/AERO47225.2020.9172374.

[138] Humboldt C. Mandell. Cost-estimating relationships for space programs. 1992.

[139] Cheng Juius Wang Ming-Wei Lu. Weibull Data Analysis with Few or no Failures, 2008.

[140] Gregory K. Mislick and Daniel A. Nussbaum. *Cost Estimation Methods and Tools*. Wiley, 2015. ISBN 9781118536131.

[141] NASA. Commercial Market Assessment for Crew and Cargo Systems. Technical report, NASA, 2011.

[142] NASA. *NASA Cost Estimating Handbook*, chapter Appendix G. NASA, 2015.

[143] NASA Associate Deputy Administrator for Policy. Falcon 9 Launch Vehicle NAFCOM Cost Estimates. Presentation, NASA, 2011.

[144] Wayne Nelson. Weibull Analysis of Reliability Data with Few or No Failures. 2018. doi: 10.1080/00224065.1985.11978953.

[145] NASA Engineering Cost Office. Crew and space transportation systems cost model (casts) user's guide.

[146] David H Olwell, Reliasoft Professional, and Services Tucson. Warranty Calculations For Missiles With Only Current-status Data , Using Bayesian Methods. 2001.

[147] P. Pempie and C. Aussilhou. Liquid Rocket Engine Test Plan and Reliability Estimation. *Japan Society for Aeronautical and Space sciences and ISTS*, ISTS 2002-a-39(January 2002):203–208, 2002.

[148] Massimo Panicucci. Expert judgment consultation - personal communication.

[149] Paolo Martino. Costs and risks analysis tool for conceptual launch vehicle MDO. 2010.

[150] R C Parkinson. The Hidden Costs of Reliability and Failure in Launch Systems. 44(99):419–424, 1999.

[151] P.Caisso, P. Brossel, T. Excoffon, M. Illig, and T. Margat. Development Status of the Vulcain 2 Engine. (July), 2001.

[152] Michel Van Pelt. Expert judgment consultation - personal communication.

[153] Pascal Pempie, Hilda Vernin, and S. A. Ligeron. Liquid rocket engine test plan comparison. *37th Joint Propulsion Conference and Exhibit*, (December), 2001. doi: 10.2514/6.2001-3256.

[154] L Pepermans. Reusable Rocket Upper Stage. 2019.

[155] Lars Pepermans. The Reusable Upper Stage. 2018.

[156] Roger Pielke and Radford Byerly. Shuttle programme lifetime cost. *Nature*, 472(7341):38, 2011. ISSN 1476-4687. doi: 10.1038/472038d. URL https://doi.org/10.1038/472038d.

[157] D. J. Rasky, F. S. Milos, and T. H. Squire. Thermal protection system materials and costs for future reusable launch vehicles. *Journal of Spacecraft and Rockets*, 38(2):294–296, 2001. ISSN 15336794. doi: 10.2514/2.3686.

[158] Daniel J. Rasky, R. Bruce Pittman, and Mark E. Newfield. The reusable launch vehicle challenge. *Collection of Technical Papers - Space 2006 Conference*, 1(September):71–81, 2006. doi: 10.2514/6.2006-7208.

[159] Vidya Sagar Reddy. The SpaceX Effect. *New Space*, 6(2):125–134, 2018. ISSN 21680264. doi: 10.1089/space.2017.0032.

[160] Georg Reibold. Esa cost engineering handbook. .

[161] Georg Reibold. Spice - standard parametric information for cost engineering cost estimating reference manual. .

[162] Reinbold. Falcon 9 Cost Estimate with SPICE. 2016.

[163] John A. Rice. Mathematical Statistics and Data Analysis, 1995.

[164] Reuben R Rohrschneider. Development of a Mass Estimating Relationship Database for Launch Vehicle Conceptual Design. 2002.

[165] M D Rozemeijer. Launch Vehicle First Stage Reusability. 2020.

[166] S Saubadine J M Ruault. EUCASS 2019 July 1 to 4 th , MADRID , SPAIN. 62:1–15, 2019. doi: 10.13009/EUCASS2019-639.

[167] Christopher Rush and Rajkumar Roy. Analysis of cost estimating processes used within a concurrent engineering environment throughout a product life cycle. 44(0), 2002.

[168] Fayssal M Safie, Marshall Space, Steven K Broussard, and Marshall Space. Reliability and Maintainability – Key To Launch Vehicle Affordability.

[169] M Snijders. Cost effectiveness of the first stage recovery of a small satellite launcher. 2017.

[170] Khanh Le Son, Laurent Saintis, Tristan Le-Fevre, and Michal Kurela. Reliability growth assessment under several operating stress conditions. *Proceedings - Annual Reliability and Maintainability Symposium*, 2015-May:2–7, 2015. ISSN 0149144X. doi: 10.1109/RAMS.2015.7105060.

[171] SpaceX. Falcon user's guide. (August), 2020. URL https://www.spacex.com/media/Falcon_Users_Guide_082020.pdf.

[172] Sippel Martin Stappert, Sven and Madalin Simioana. Technical report on different rlv return modes' performances.

[173] Mark G. Stewart. Reliability-based assessment of ageing bridges using risk ranking nd life cycle decision analyses.

[174] Paul V Tartabini, James R Beaty, Roger A Lepsch, and Michael G Gilbert. Payload Performance Analysis for a Reusable Two-Stage-to-Orbit Vehicle. (May 2015), 2019.

[175] ECOS Team. Ecos 5 training - self training manual. technical report, european space agency. .

[176] NASA Independent Review Team. Spacex crs-7 accident investigation report. .

[177] Kong Fah Tee. Reliability based life cycle cost optimization for underground pipeline networks.

[178] Thiokol Propulsion. Launch vehicle historical reliability. 1998. doi: 10.2514/6.1998-3979.

[179] O Trivailo, M Sippel, and Y A Şekercioğlu. Review of hardware cost estimation methods, models and tools applied to early phases of space mission planning. *Progress in Aerospace Sciences*, 53:1–17, 2012. ISSN 0376-0421. doi: http://dx.doi.org/10.1016/j.paerosci.2012.02.001. URL http://www.sciencedirect.com/science/article/pii/S0376042112000140.

[180] US Army Materiel Systems Analysis Activity. AMSAA Technical Report No. 138 REliability Analysis for Complex Repairable Systems. 1976.

[181] J. M. van Noortwijk. A survey of the application of gamma processes in maintenance. *Reliability Engineering and System Safety*, 94(1):2–21, 2009. ISSN 09518320. doi: 10.1016/j.ress.2007.03.019.

[182] Fred E Wagner. EXPECTED VALUE ANALYSIS FOR AN UNMANNED EXPENDABLE LAUNCH VEHICLE PAYLOAD ESCAPE SYSTEM. 1987.

[183] Matthew Weinzierl. Space, the final economic frontier. *Journal of Economic Perspectives*, 32(2):173–192, 2018. ISSN 08953309. doi: 10.1257/jep.32.2.173.

[184] James R. Wertz. Economic model of reusable vs. expendable launch vehicles. .

[185] James R. Wertz. Responsive launch vehicle cost model. .

[186] Robert Williams and Kevin Walsh. Covering the Increased Liability of New Launch Markets. pages 1–9, 2016.

[187] Qin Xu, Peter Hollingsworth, and Katharine Smith. Launch cost analysis and optimization based on analysis of space system characteristics. *Transactions of the Japan Society for Aeronautical and Space Sciences*, 62(4):175–183, 2019. ISSN 05493811. doi: 10.2322/tjsass.62.175.

[188] Shigeru Yamada and Shunji Osaki. Software Reliability Growth Modeling: Models and Applications. *IEEE Transactions on Software Engineering*, SE-11(12):1431–1437, 1985. ISSN 00985589. doi: 10.1109/TSE.1985.232179.

[189] J. N . Yang. Inspection frequency optimization for aircraft structures based on reliability analysis.

[190] Anthony Young. *The Saturn V F-1 Engine - Powering Apollo Into History*. 2008. ISBN 9780387096292.

[191] David A. Young. An innovative methodology for allocating reliability and cost in a lunar exploration architecture.

[192] B T C Zandbergen. AE1222-II : Aerospace Design Systems Engineering Elements I Part : Launcher design. 2021.

[193] Barry Zandbergen. Simple mass and size estimation relationships of pump fed rocket engines. .

[194] Barry Zandbergen. Ae1222-ii: Aerospace design systems engineering elements i. .

[195] Franziska Zilker. Aerothermal analysis of re-usable first stage during rocket retro-propulsion. (January):35–61, 2018.

[196] Mathew R Zwack. Contrast : a Conceptual Reliability Growth Approach for Comparison of Launch Vehicle Architectures. (December), 2014.

[197] Nedim Sözbir Şenol Gülgönül. Selection of geostationary satellite launch vehicle using expected value analysis.

[198] şenol Gülgönül and Nedim Sözbir. Selection of Geostationary Satellite Launch Vehicle Using Expected Value Analysis. *Sakarya University Journal of Science*, 22(5):1–1, 2018. ISSN 1301-4048. doi: 10.16984/saufenbilder.346597.

Appendices

A

Appendix A: Inflation Table

The value of money changes every year based on inflation. The CPI provides a measure of inflation. Values from Organisation for Economic Co-operation and Development (OECD) are provided here for the United States, European Union and OECD, from 2010 up to 2021.

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
USA	92	94.9	96.9	98.3	99.9	100	101.3	103.4	105.9	107.9	109.2	114.3
EU	93	95.7	98.2	99.5	99.9	100	100.2	101.7	103.6	105	105.8	108.8
OCDE	91.4	94	96.1	97.6	99.3	100	101.2	103.5	106.2	108.4	109.9	114.3

Table A.1: Consumer Price Index table, reference year 2015 [1].

References

[1] Consumer Price Indices (CPIs) - Complete Database. https://stats.oecd.org/index.aspx? DataSetCode=PRICES_CPI#. Accessed: 12/05/2022.

B

Appendix B: Currency Conversion Table

Data for conversion between Euro and Dollar was taken from the European Central Bank.

2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
1.3257	1.3920	1.2848	1.3281	1.3285	1.1095	1.1069	1.1297	1.1810	1.1195	1.1422	1.1827

Table B.1: Euro to Dollar conversion table [1].

References

[1] Euro foreign exchange reference rates. https://www.ecb.europa.eu/stats/policy_and_ exchange_rates/euro_refere
graph-usd.en.html. Accessed: 12/05/2022.

C

Appendix C: Work-Year Costs Table

Values for Work-year costs are taken from TRANSOCST [129]. Values for years after 2014 are extrapolated linearly from the known data from previous years.

YEAR	USD	EUR
1990	156,200	139,650
1991	162,500	145,900
1992	168,200	151,800
1993	172,900	156,800
1994	177,200	160,800
1995	182,000	167,300
1996	186,900	172,500
1997	191,600	177,650
1998	197,300	181,900
1999	203,000	186,300
2000	208,700	190,750
2001	214,500	195,900
2002	222,600	201,200
2003	230,400	207,000
2004	240,600	212,800
2005	250,200	219,200
2006	259,200	226,300
2007	268,800	234,800
2008	278,200	243,600
2009	286,600	252,700
2010	296,000	261,000
2011	303,400	268,800
2012	312,000	275,500
2013	320,000	285,000
2014	328,700	292,400
2015	327,711	288,589
2016	335,019	294,765
2017	342,327	300,941
2018	349,635	307,117
2019	356,943	313,293
2020	364,251	319,468
2021	371,559	325,644

Table C.1: Work-year Costs [1].

References

[1] Dietrich E Koelle. *Handbook of Cost Engineering and Design of Space Transportation Systems*. Number TCS-TR-200. TransCostSystems (TCS), Ottobrunn Germany, 8.2 edition, 2013.