UTILITY OPTIMAL DECISION MAKING WHEN RESPONDING TO NO FAULT FOUND EVENTS

by

Archana Ravindran

A Thesis

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Master of Science in Aeronautics and Astronautics

School of Aeronautics and Astronautics

West Lafayette, Indiana

August 2020

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL

Dr. Karen Marais, Chair

School of Aeronautics and Astronautics

Dr. William A. Crossley

School of Aeronautics and Astronautics

Dr. C. Robert Kenley

School of Aeronautics and Astronautics

Approved by:

Dr. Gregory Blaisdell

Dedicated to my parents

ACKNOWLEDGMENTS

I would like to thank my advisor Prof. Marais, for her guidance and encouragement. Her knowledge, experience and insights were elemental in this research work. She takes keen interest in her students' success and motivates them through the struggles of graduate life.

My friends in the research lab, for listening to countless iterations of my work, giving your valuable inputs and making the office feel like home.

My research committee members, Prof. Crossley and Prof. Kenley for their patience, understanding and ideas in taking this research forward.

Luna Burgos Moreno, who started this work before me and laid its foundation.

Varun Sudarsanan for being my personal cheerleader and soundboard for all my research ideas and roadblocks.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ATA	=	Air Transport Association
DM	=	Decision Maker
LRU	=	Line Replaceable Unit
MRO	=	Maintenance, Repair and Overhaul
NFF	=	No Fault Found
NPV	=	Net Present Value
OEM	=	Original Equipment Manufacturer

ABSTRACT

No Fault Founds (NFFs) are an expensive problem faced by the airline industry. The underlying cause of NFFs are a major focus of research work in the field, but the dearth of consistent data is a roadblock faced by many decision makers. An important risk factor identified is the occurrence rate of NFFs.

This research work aims to help decision makers in the Airline Maintenance, Repair and Overhaul teams, when faced with recurring NFFs, to make a choice based on value derived from the system and risk preference of the decision maker under uncertainty. The value of the aircraft fleet is laid out using Net Present Value at every decision point along the system life cycle while accounting for the uncertainty in the failure rate information. Two extreme decisions are considered for the decision maker to choose between: rebooting the system every time a failure occurs and results in an NFF which allows for it to recur while reducing uncertainty of the failure rate; or eliminating the failure mode which assumes that the failure does not recur and therefore completely removes the uncertainty. Both decisions have their associated uncertain costs that affect the NPV calculated. We use a Monte Carlo approach to estimate the expected profit from deciding to eliminate the failure mode. We make use of Expected Utility Theory to account for the risk preference of a decision maker under uncertainty and build an Expected Utility Maximizing decision framework.

To conclude we give some guidance to interpret the results and understand what factors influence the optimal decision. We conclude that not accounting for uncertainty in estimating a failure rate for the future along with uncertainty in NFF costs can lead to an undesirable decision. If the decision maker waits too long to gather more information and reduce uncertainty, then rebooting the system for the remaining life could be more worthwhile than spending the large amount of money to Eliminate a failure mode. Finally, we conclude that, despite uncertainties in information of occurrence rates and costs of NFFs, an Expected Utility maximizing decision between the two options considered – Reboot and Eliminate – is possible given the available information.

1. INTRODUCTION

The Air Transport Association (ATA) has reported that 4,500 No Fault Found (NFF) events cost ATA member airlines $100 million annually, while also causing numerous flight delays and cancellations (Beniaminy & Joseph, 2002). It is estimated that 30-50% of the avionics failures in the aviation industry are NFF events (Khan et al., 2014). Each year, commercial airlines in the United States spend about $185,000/aircraft on unsuccessful attempts at replicating reported avionics failures (Werner, 2015). These estimates are likely conservative since there are other auxiliary costs of spare parts, human resource, and revenue loss due to downtime which can inflate the cost associated with NFFs.

There are different definitions for what constitutes an NFF event. In the ARINC 672 report, it is described as "Removal of equipment from service for reasons that cannot be verified by the maintenance process (shop or elsewhere)" (ARINC 672, 2008). An NFF event implies that a failure (fault) either occurred or was reported to have occurred during a system's use, but upon subsequent investigation there was either no evidence of the failure (e.g., a burnt-out circuit), or the failure could not be replicated (Qi et al., 2008). Such NFF events are variously referred to as trouble-not-identified (TNI), cannot duplicate (CND), no-trouble-found (NTF), and retest OK (RTOK).

NFFs predominantly occur in electronic devices (Khan et al., 2014) and are consequently encountered in all industries including the automotive, avionics, telecommunications, computer, and consumer industries. Field returns that result in NFFs can be debilitating for the OEMs. The manufacturer may receive an unfavorable reputation, and replacements can be costly if the product is within the warranty period. In the 1980s, Ford's engine ignition module saw a four-fold increase beyond the projected warranty returns for a 30-month period (Thomas et al., 2012). These returns were categorized as NFFs since the modules that were returned were fully functional when re-tested. Ford later identified many failure modes that could lead to this component's intermittent failure. Ford was taken to trial and was mandated to recall the module. Ford was required to extend the warranty of the module and cover all the replacement costs. Ford also agreed to contribute USD 5 million academic research towards automotive safety.

When a failure occurs, there are multiple maintenance strategies that can be applied to restore the system. The commonly employed strategy is preventive maintenance, where there is a predetermined maintenance schedule. For example, according to airworthiness regulations, an aircraft has to go through A-check maintenance depending on flight hours, cycles and calendar months (Kinnison & Siddiqui, 2013). This type of maintenance arises from an understanding of how the product functions and the wear and tear of its parts over time. Another type of maintenance is predictive maintenance, where the product has continuous health monitoring and/or sensors installed. The data collected from the sensors allow the maintenance crew to determine the response strategy.

The approaches mentioned above are proactive in nature. However, maintenance is reactive in nature when unexpected faults occur, and the failed equipment has to be repaired or replaced after the failure occurs (Swanson, 2001). In the safety and cost critical airline industry, it is important to minimize the downtime while carrying out reactive maintenance. Diagnosing and repairing an avionics fault can be time consuming due to the complexity of the system. The concept of Line Replaceable Units (LRU) has emerged to achieve a quick turnaround time while adhering to the safety standards.

For avionics, a typical reactive maintenance process of fault identification and repair starts with an operator (e.g., pilot) reporting an error. The maintenance technician replaces the component with an LRU and sends the faulty component to the repair shop. If the repair shop cannot replicate the reported fault, the component is tagged as an NFF and put back into stock.

An NFF does not imply that the fault does not exist. It is a failure to detect the error and therefore is a deficiency in the fault-finding process. Examples of such deficiencies include the operator/maintenance technician not fully understanding how the part functions, incorrect or ambiguous fault diagnosis manual, or a lack of resources in the repair shop to detect the fault.

NFFs can be dangerous. In September 2010, during final approach, the crew of a Bombardier Dash 8 Q400 was so distracted by a non-functioning flight display, that they inadvertently disabled the autopilot and would likely have hit the ground, had the ground proximity warning not been activated (Werner, 2015). The same underlying problem with the input/output processor had occurred on several other flights, but each time technicians were unable to replicate the problem on the ground.

Even when there is no direct impact on safety, NFFs can still have significant costs, both tangible (e.g., decreased reliability and availability, the costs of attempting to replicate and correct the fault, and warranty costs) and intangible (e.g., customer perception of inadequate quality).

In general, when failures occur, organizations are faced with a range of choices. At one extreme, they choose to act and remove the particular failure mode. At the other extreme, they can choose to repair or replace failed or faulty components. The most appropriate choice depends on several factors, including whether the failed component is safety critical, the impact the failure has on system performance, the cost and feasibility of identifying and removing the failure mode, and the cost and feasibility of (perhaps repeatedly) repairing failed components. There is also a tradeoff along the time dimension in making a choice—act early with limited information, possibly investing in preventing a failure that would have been infrequent, or, wait for more information, possibly incurring frequent costly failures.

In the case of NFFs, these extreme choices translate into identifying and eliminating the source of the NFF, or to "accept" the NFF and only do what is necessary to recover from it (e.g., reboot a computer). What is the best choice in a given situation? Here, we present a decision framework based on Net Present Value and Expected Utility to aid in such decisions.

1.1 Literature review

The NFF phenomenon is prevalent across many industries such as the automotive, consumer electronics, aviation and space sectors. Research efforts addressing the NFF phenomenon have gained traction in the aerospace sector in the past decade (Khan et al., 2014).

Söderholm (2007) says that "Traditionally, when an NFF event was encountered, the conclusion was that there was no fault present in the system". The NFF category would not contribute to the failure statistics in many companies (Qi et al., 2008). But recently, this attitude towards NFFs is seeing a shift. Qi et al. suggest that intermittent faults are a common cause of NFF events and advocate for NFFs to be treated and investigated as failures to identify their root causes. They state that “organizations may not understand the need or have little incentive to uncover the root cause of the problem encountered by the user”.

There have been efforts to classify NFFs (Khan et al., 2014) and describe the NFF phenomenon at various levels of the maintenance process (Khan et al., 2017; Söderholm, 2007). In a 2007 review paper, Söderholm et al. describe the NFF phenomenon and provide possible improvements to prevent NFFs and reduce the consequences of NFFs on the lifecycle stages, availability of spare units and key stakeholders of the system. They state that an important risk aspect of the NFF is its frequency of occurrence. It is estimated that, in the aerospace industry, about 50 percent of LRUs that are removed and replaced during operation and maintenance are classified as NFF events (Beniaminy & Joseph, 2002; James et al., 2003; Sudolsky, 1998).

Khan et al., (2017) focus on decision making at the operational, tactical and strategic level when encountered with NFFs. Compared to standardized maintenance processes, resolving NFFs require more control and escalation of decision making due to the high-pressure situations at the operational level. He goes on to suggest that a Petri net modelling of the decision process can help in mathematically modelling a unified framework that incorporates different scenarios and alternatives when faced with NFF issues. In a previous review paper, Khan (2014) implores that there needs to be a standardized taxonomy for the NFF phenomenon, the lack of which adds to the problem of sparse knowledge in the field of NFFs. They state that there is a lack of serious research in the direction of understanding the relationship of type of equipment, its usage and complexity to the rate of occurrence of NFFs.

The body of literature mainly focuses on understanding why NFFs occur and what can be done to reduce their occurrence and mitigate their effect. Erkoyuncu (2016) takes a step towards narrowing down the major NFF cost drivers across the supply chain and build a framework to estimate the cost of NFF events. They interviewed industry engineers to understand the process followed in the industry to deal with NFFs and found that it varies from organization to organization. The study provides cost of an NFF but does not account for the uncertainty in the rate of NFF occurrence and the effect of the different decision options available to the decision maker on the overall value provided by a system.

In the current work we attempt to provide a framework for optimal decision making under uncertainty of NFF occurrence rates and costs. The discounted cash flow technique of Net Present Value (NPV) and Expected Utility Theory are used as the basis for the framework for the decision maker to decide between two possible choices.

1.2 Aircraft Maintenance Process

In order to model the failure rates and occurrence of NFF events along with the operating and maintenance costs, it is important to understand the maintenance process undertaken for an aircraft. We can then establish what process would be applicable in case of deeming a failure to be an NFF and the different costs that go into bringing the system back to a functional state.

Material for this section is heavily borrowed from the books, *Aviation Maintenance Management* by Harry A. Kinnison and *Leveraging Information Technology for Optimal Aircraft Maintenance, Repair and Overhaul (MRO)* by Anant Sahay. The reader is encouraged to refer to these books for more information on this subject.

1.2.1 The Maintenance Steering Group (MSG) process

A study by United Airlines showed that only 11% of items included in the study would benefit from having scheduled maintenance checks. The other 89% would require other maintenance programs.

A Maintenance Steering Group (MSG), formed by Airlines for America, has developed a handbook of requirements for scheduled maintenance procedures which has been revised over time and is now called MSG-3 (Air Transport Association of America [ATA], 2002). In 1968, representatives from various airlines formed a Maintenance Steering Group (MSG) and developed Handbook MSG-1, "Maintenance Evaluation and Program Development," which included decision logic and inter-airline/manufacturer procedures for scheduled maintenance for the new Boeing 747 aircraft. This document was later modified to be universally applicable to new aircraft. The updated decision-logic, MSG-2, was used to develop scheduled maintenance for aircraft of the 1970s. The MSG-2 process classified each unit (system or component) to one of the assigned maintenance processes. These processes were the Hard Time (HT), On-condition (OC) and

Condition monitoring (CM) (Kinnison & Siddiqui, 2013). Hard Time is preventive maintenance technique which requires component to be removed at a predetermined interval of operating time. On-condition is when a component's remaining serviceability is checked periodically. Condition monitoring is when failure rates and deterioration of the component are monitored for maintenance planning of the component.

In 1980, the MSG-2 process was modified to adopt a task-oriented maintenance approach. The updated process, called MSG-3, is a task-oriented approach to decision logic. It is developed by the Air Transport Association and identifies scheduled maintenance tasks which allow for the reliability of the system to be maintained. The MSG-3 is a top-down approach and considers how the failure of a component affects the aircraft operation (ATA, 2002). It has two-levels of classification of components. The level I analysis is to assign the failure into two basic categories: evident and hidden to operating crew. These are further split into safety related and operationally related failures. In the level II analysis, the maintenance tasks required are determined using a question flow chart.

Working groups for different systems on the aircraft will receive information about the system. This information could be the theory or operation and its modes, failure modes of each operation and any data collected (like failure rates). The working groups receive training which acts as a refresher, if a similar system exists already, or allows them to fully understand the failure modes of the new or updated system. The manufacturer is responsible for the training and for providing this information to the working groups. Once the group is sufficiently informed, they begin running through the logic diagrams to determine the best maintenance approach for each component, each operational mode and each failure mode in their assigned system. The group, using their knowledge and judgement, also determines at what intervals the maintenance tasks need to be performed.

1.2.2 Procedures to combat in-service / unscheduled failures

The MSG-3 approach to maintenance is used to avoid in-service failures. Equipment redundancy, Line Replaceable Units (LRUs) and Minimum Equipment List (MEL) are some management techniques used to allow failures while in-service without affecting safe and timely operation (Kinnison & Siddiqui, 2013). LRUs are components that are designed such that, parts which experience failures commonly, can be quickly removed and replaced on the vehicle. Modern avionics are modular with a set of LRUs that are identical to the ones in operation and are easy to access, remove and replace. An LRU can have its own Built-in test equipment (BITE) and can consist of several Shop Replaceable Units (SRUs) (Raza & Ulanskyi, 2016). The BITE provides continuous testing output to indicate the health of the component. The LRUs are run-to-failure (which can be either permanent failures or intermittent), and these failures are recorded while in operation (during flight) or on the ground. LRUs that are removed from the aircraft are re-tested and repaired at the maintenance shop or are sent to the manufacturer to repair. An Automatic Test Equipment (ATE) is necessary for retesting LRUs that are dismantled and to detect a failed SRU. Redundancy is built into avionics systems on the aircraft to provide safe operations and sufficient LRUs are kept in-stock to maintain operations.

The MEL allows for the vehicle to be in service even if some parts are inoperative, provided that it does not affect safety and operation of the flight. The manufacturer of the aircraft provides the Master Minimum Equipment List (MMEL) and the airline tailors it to suit its needs to make the MEL. Many MEL items have associated redundancy.

The described maintenance strategies can be inefficient when NFFs events occur because testing equipment might not identify the intermittent failures and they may re-occur in the future. When the failure rates of components are not predictable and there is no scheduled maintenance procedure in place for these components, an MRO usually uses a reliability program.

1.2.3 Reliability program within the Maintenance and Engineering (M&E) area

The primary approach to a reliability program is to specifically address maintenance problems (even the ones that do not cause delays). The main functions of the reliability program are: 1. Monitor the performance of the vehicles and their equipment and call attention to any need for corrective action; 2. Monitor the effectiveness of said corrective actions; 3. Provide data to justify adjusting the maintenance intervals or maintenance program procedures whenever those actions are appropriate. The basic tasks taken up by the reliability program to feed into the functions above are:

a. Data collection – this includes flight time and cycles (since most failure/removal rates are based on flight hours or cycles), delays over 15 minutes (specifically the ones due to maintenance), unscheduled component removals (if the rate is not acceptable, investigation and corrective actions are undertaken), pilot reports and logbook reports, component failures in shop maintenance, maintenance check findings.

b. Problem area alerting – Standard event rates are set based on past performance. An Upper Control Limit (UCL) is used as an alert limit. The UCL above the mean value by 1 to 3σ (standard deviations).

c. Data display and analysis – the reliability department does a preliminary analysis of data before alerting the engineering department to establish its validity.

d. Corrective action – the engineering team will investigate the problem and determine required redesign or required change in maintenance procedures. The team then issues an engineering order for implementation of the required action. This action needs to be approved by the Maintenance Program Review Board (MPRB) before being undertaken by the reliability and maintenance teams.

e. Follow-up analysis – the reliability department continues to monitor the effect of the corrective action on on-alert items. The effectiveness of the corrective actions is reflected in decreased event rates.

f. Data reporting – the reliability report is issued monthly and contains information on on-alert items and items under follow-up investigation. The report is organized by fleet, this means each aircraft type is addressed separately in the report.

As airlines started to grow and the volume of work undertaken by the maintenance and engineering teams increased, airlines identified aircraft maintenance as a non-core function (Sahay, 2012). Airlines started creating a separate independently operating maintenance engineering unit, either wholly owned subsidiaries or independent companies. These new entities, called Maintenance, Repair and Overhaul (MRO) organizations, either operate as profit centers for the mother airline or independent companies and add to their revenue by providing services to low cost carriers (LCC). The FAA in its documents generally assumes that maintenance of a commercial aircraft is done by the department or a division of an airline. However, the current trend is for MRO organizations to be set up either independently or as joint ventures with airlines and OEM.

In this research work, we assume that the airline operates and bears the cost for maintenance of its fleet. Therefore, the operating cost, net profit, cost of maintenance when NFF events are identified (such as cost of spare LRU units, redundant systems and movement of the LRU and SRU unit from on-ground service to repair shops and back into stock pile), are attributed to and borne by the airline. The collection of data for failure/removal rates, flight hours and hours of operation and downtime are also assumed to be tasks carried out by the M&E department of the airline.

In the next section, we discuss the building blocks of the decision-making framework. The failure model for NFFs and their uncertainty, Net Present Value and Expected Utility Theory are discussed along with how the DM can use these to make the optimal choice. This section is followed by the results section in which we elaborate on the behavior of NPV and Expected Utility, their sensitivity to different inputs and develop a concept of threshold failure rate to help the DM with decision making. Section 4 concludes the work and summarizes the contributions of this research.

2. THEORY: A VALUE AND UTILITY PERSPECTIVE OF A DECISION

This section discusses the modelling approach and methodology of building the framework for optimal decision making with uncertain failure rates.

We first model the fleet as a revenue generating system. We then see how failures affect the cash flow of this system. We make some assumptions that help build the decision framework. Using these assumptions, we lay out how failures are modeled for the two decisions we consider – Eliminate and Reboot – and how NPV is calculated for both. Following this, we explain how NPV can help decide between the two options, but uncertainty can lead to an undesirable choice.

We consider three different sources of uncertainty in our work. The first is the uncertainty in failure rate due to imperfect knowledge of the NFFs characteristics. The other two uncertainties are costs required to either Reboot the system or Eliminate a failure mode. The three uncertainties are considered together to calculate Utility of the Eliminate option. Here we consider the risk preference of the DM in an attempt quantify the risk in choosing to Eliminate at a given decision point under the given uncertainty.

2.1 Prelude: Systems as Revenue Generators that experience failures

The approach builds on previous work (Marais, 2013; Marais & Saleh, 2009; Saleh & Marais, 2006), and is based on viewing the system as a revenue generating unit (Figure 1). Value is then the difference between revenue and cost. For example, the revenue seat miles provided by a commercial passenger aircraft can be used to calculate its value, while a telecommunications satellites has transmitters which provide bandwidth and the value of this satellite can be calculated using this bandwidth.

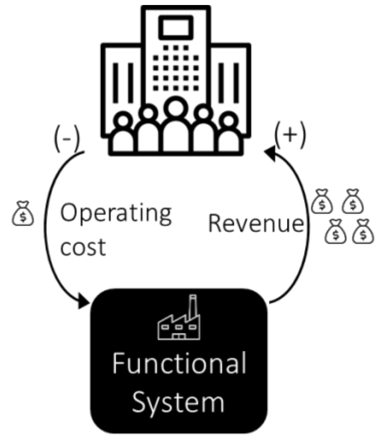

Figure 1: System is a revenue generating artifact with operation cost being the cost put in to run the system

Considering an aircraft as an example of a functional system, we define the operating cost as the cost put in by an airline to run the aircraft. It includes fuel cost, ownership cost, salary for the crew, and maintenance cost. Revenue is generated from ticket sales.

We look at value of a system from the perspective of net cash flow through the system for a given time interval, t:

$$Value(t) = Revenue(t) - Operating\ cost(t) \tag{1}$$

This values for all time intervals are summed over the full lifetime of the system and discounted to get the Net Present Value:

$$NPV = \sum_{t=0}^{End\ of\ life} \frac{Value(t)}{(1+r)^t} \tag{2}$$

Where r is the discount rate to account for the time value of money.

When a failure occurs, the system goes from a functional state to a reduced functional state and requires some maintenance activity to bring it back to its functional state (Figure 2). Failures can have several impacts on profit, including direct costs like cost of new components and the downtime associated with the failure and subsequent repair, and indirect costs like reduced customer satisfaction. For example, a pilot reports a glitchy communications radio during pre-flight inspection. The flight may have to be delayed while the event is logged and inspected. The airline may incur a loss in revenue due to the delay, additional operating costs (e.g. crews working extra time), and manpower and material costs used to investigate the fault. The maintenance technician may choose to replace the component with an LRU and send the faulty component to the shop for testing and repair. The maintenance shop may fail to diagnose the fault and mark the event as an NFF and place the component back into the LRU stock pile. This chain of events will have a cost associated with it (for example the cost of logging the failure, cost of extra human resource that might have to be deployed, cost of testing and extra LRUs required) and is incurred each time an NFF event occurs.

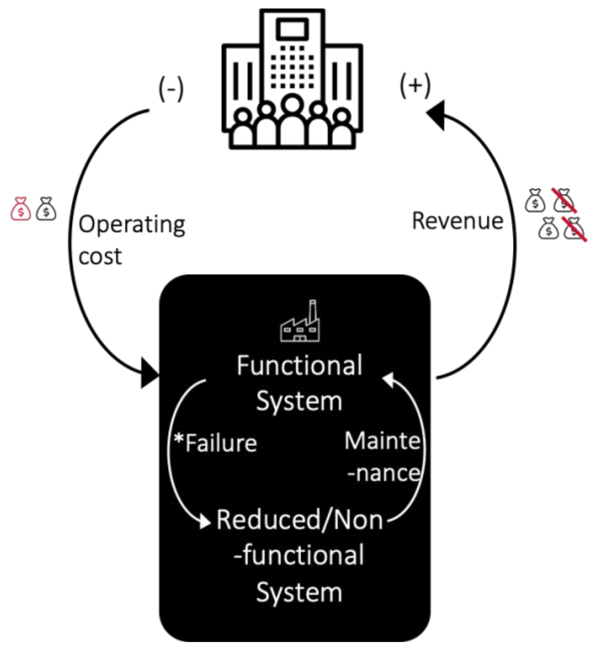

Figure 2: Failures require additional maintenance cost to bring the system back to its original functional state

Failures tagged as NFFs can recur if left unchecked. If the NFF is associated with a component that is cheap, difficult to inspect, easy to replace, and does not have a safety critical function (e.g., a lightbulb), replacement on failure may be the best option. If a component's failure entails significant downtime and cost (e.g., an aircraft empennage), but the root causes (e.g., metal fatigue) of the failure are difficult or impossible to eliminate, inspection and preventive maintenance may be the best option. And if a component's failure entails significant and possibly frequent downtime and cost (e.g., a software bug), but it is possible that the root causes (e.g., incorrect specification) of the failure can be found, eliminating the failure mode may be the best option.

In this research we attempt to provide the Decision Maker (DM) with discounted cash flow analysis of the system for different the different maintenance strategies they choose. We consider a single type of NFF that occurs in a particular component and assume it has a single failure mode. We model the effect of recurring failures due to this NFF on the cash flow through 1) repair costs that are added on to the operating cost and; 2) downtime due to failures that reduces the revenue generated. We assume that all other scheduled and unscheduled maintenance has their own cost and they are all accounted for in the operating cost. We also assume that when a failure is reported, the aircraft experiences downtime and is no longer in a functional state and needs replacement of the LRU to become functional again.

2.2 A model of Net Present Value in the presence of NFFs

2.2.1 Assumptions in building the model

In developing our model, we make some assumptions that help building the model of the cash flow of the aircraft system and keep the focus on the main argument of this work:

Assumptions for constraining the type of NFF

1. We consider only non-safety-critical failures.
2. Each NFF corresponds to a single failure mode.
3. We assume that there is no historic data of NFF failure rate available for use at the start of the study.
4. The component for which the NFF problem is studied is identical or similar across the fleet of the airline. This assumption allows the use of a single failure rate to analyze the entire fleet.

Assumption for the failure model and cash flow model

1. Failures can be modelled as stochastic process. In particular, we assume that NFFs can be modeled as following a Homogenous Poisson Process (Montgomery, 2005). This assumption makes some of the equations easier to present but can be changed easily and does not affect the framework.

2. The airline fleet is homogenous. This assumption allows for costs to be constant across the fleet.
3. Revenue, operating, repair, reboot, eliminate costs remain same for all time intervals, that is, we do not account for external effects on costs such as market forces.
4. We assume the downtime and the cost during downtime are same for all failures. The cost to reboot and to eliminate can be different from each other.
5. We assume that failures are reported and investigated as soon as they occur, and therefore, we can conflate the initial failure and subsequent NFF in our definition of downtime, reboot cost, and loss of revenue.
6. We assume that the impact of all failures other than the particular mode of NFF failures investigated, is reflected in the nominal operating cost and revenue.
7. Each month has 30 days, and hence each year has 360 days.
8. The problem time horizon is finite and set to 30 years. This lifetime is in line with lifetimes of complex engineering systems like aircraft (BTS, 2018).

Assumptions for the decision model

9. We assume the time interval between two decision making points is one month. This assumption is in line with airline Maintenance Review Board having monthly meetings to discuss status of on-alert items (Kinnison & Siddiqui, 2013). It is assumed that the decisions to eliminate the failure mode or continue to reboot the system are made at the end of a month.
10. Decisions to reboot or eliminate are made at the end of each time interval. To simplify calculations, we also assume that the total downtime in any given month, associated with either rebooting or eliminating the failure mode, is less than or equal to one month, i.e.:

    $$\lambda_i \cdot t_R \leq \Delta T \; and \; t_E \leq \Delta T$$

    Where λ_i is the failure rate for time interval i, ΔT is the time interval between two decision points, t_R is the downtime when rebooting and t_E is the downtime when eliminating.
11. If chosen, the eliminate option is always successful in eliminating the failure mode.
12. The reboot option, when chosen, does not provide any knowledge about the failure mode except its frequency. This would imply that all the 'learning' we do about the failure mode is assumed to occur if and when eliminate option is chosen.

2.2.2 Modeling failure rate

To account for cost due to NFF each time a failure is reported, we need to record the number of failures every month and have an estimate of failure rates for future months. In quality control, typically a Poisson distribution is used to model the number of failures in a product unit per unit time. The probability distribution function for the Homogeneous Poisson Process is given by:

$$P(X = x) = e^{-\lambda} \frac{\lambda^x}{x!} \; for \; x = 1, 2, 3 \ldots \tag{3}$$

Where X is the number of failure events in a given time interval and λ is the *expected number of events per time interval.*

In the absence of real recorded failure data from MRO organizations for this research, we simulate failure rates within reasonable ranges using available data.

Let us assume that the given failure mode being investigated has an underlying true failure rate, λ_{true}. Using λ_{true}, we generate a random failure rate for every month using the Poisson probability distribution. At a given decision point, j, we know the failure rates of all the past months and estimate the failure rate for the future using the Maximum Likelihood Estimate (MLE) of the mean failure rate. The MLE at decision point j, is $\hat{\lambda}_j$ and is the mean of the failure rates of all previous months (Kumar et al., 2006):

$$\hat{\lambda}_j = \frac{1}{j} \cdot \sum_{k=1}^{j} \lambda_k \tag{4}$$

Where λ_k is the number of failures across the fleet that resulted in an NFF recorded for month k. $\hat{\lambda}_j$ represents the failure rate estimated for the subsequent months from decision point, j.

The estimated failure rate will likely change as we obtain more data. Figure 3 shows an example of estimates over a one-year period with an underlying $\lambda_{true} = 15$. The estimated mean failure rates are shown with their confidence intervals. The figure also shows the simulated random sequence of actual failure rates at the end of each month.

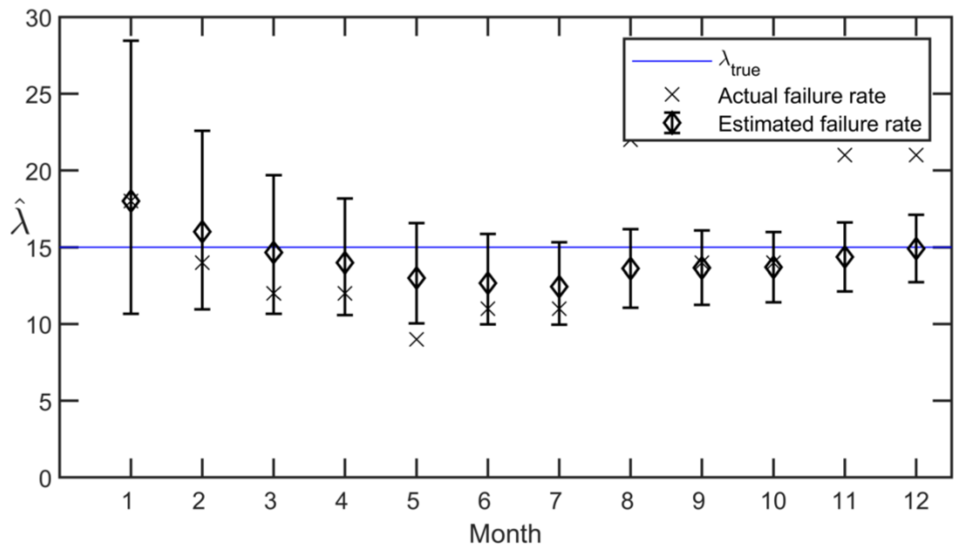

Figure 3:Failure rate estimates are updated at the end of each month using the data recorded in that month. The uncertainty interval reduces as more data is collected.

The total number of failures recorded until month j, is N_j:

$$N_j = \sum_{k=1}^{j} \lambda_k \tag{5}$$

When $N_j < \sim 100$, a two-sided confidence interval for $\hat{\lambda}_j$ is given by (Sahai & Khurshid, 1993):

$$\begin{aligned} lower\ bound &= \frac{\chi_2^{-1}\left(\frac{\alpha}{2}, 2 \cdot N_j\right)}{2 \cdot N_j / \widehat{\lambda_j}} \\ upper\ bound &= \frac{\chi_2^{-1}\left(1 - \frac{\alpha}{2}, 2 \cdot \left(N_j + 1\right)\right)}{2 \cdot N_j / \widehat{\lambda_j}} \end{aligned} \tag{6}$$

Where, $\chi_2^{-1}(p, v)$ is the inverse Chi-squared function with probability *p* and *v* degrees of freedom.

When $N_j > \sim 100$, a two-sided confidence interval for the MLE is given by (Sahai & Khurshid, 1993):

$$\begin{aligned} lower\ bound &= \frac{Z^{-1}\left(\frac{\alpha}{2}, N_j, \sqrt{N_j}\right)}{N_j / \widehat{\lambda_j}} \\ upper\ bound &= \frac{Z^{-1}\left(1 - \frac{\alpha}{2}, N_{j_j}, \sqrt{N_j}\right)}{N_j / \widehat{\lambda_j}} \end{aligned} \tag{7}$$

Where $Z^{-1}(p, \mu, \sigma)$ is the inverse normal function with probability p, mean μ, and standard deviation σ.

The Poisson distribution is a skewed distribution (unlike a normal distribution, which is symmetric). The left tail of the distribution cuts off at zero (e.g., failure rate is zero), but the right tail is long (goes to infinity). When the mean of the distribution is closer to zero, the distribution is more skewed than if the mean were far from zero. This skewness also reflects in the confidence interval – the lower bound of the confidence interval is smaller than the upper bound.

The confidence intervals for the estimated failure rates is the uncertainty in our estimates. As more data is gathered with each month, the uncertainty in the estimate reduces and we expect the estimated mean to be closer to the underlying true failure rate, λ_{true}.

2.2.3 Calculating NPV using estimated failure rate

In this research we look at two extreme decisions that the DM can take when encountered with NFFs. The first type of decision is to eliminate the failure mode by understanding the root causes of the failure and redesigning the system such that the failure does not occur again in the given failure mode. This decision is one extreme of the possible decisions because it completely removes the possibility of any failure in the future and therefore removes the uncertainty of the failure rates. The other extreme of the decisions is to continue to reboot the system until the next decision point. Rebooting a computer when it hangs or replacing a light bulb when it goes out are examples of the reboot decision scenarios. In the case of an aircraft, it means we do not take any measures to understand what is causing the failure and continue to replace the faulty unit with an LRU in stock. The Reboot and Eliminate decisions have different consequences and result in different costs incurred and therefore give different NPVs of the system.

2.2.3.1 NPV when decision is to Reboot

We illustrate our approach using an example. Assume that three NFFs were logged across the fleet for a given month. Each time a failure is reported and an LRU is replaced, the system experiences some downtime of t_R. During downtime, the system incurs additional downtime cost at a cost rate $\dot{c}_D$, (an example of these costs is the overtime pay for maintenance and ground crew). The reboot cost is c_R. We assume that rebooting always take the same amount of time/labor and is therefore constant for a given NFF. Figure 4 shows the state of the system in the 'Reboot' scenario and its corresponding costs for month j.

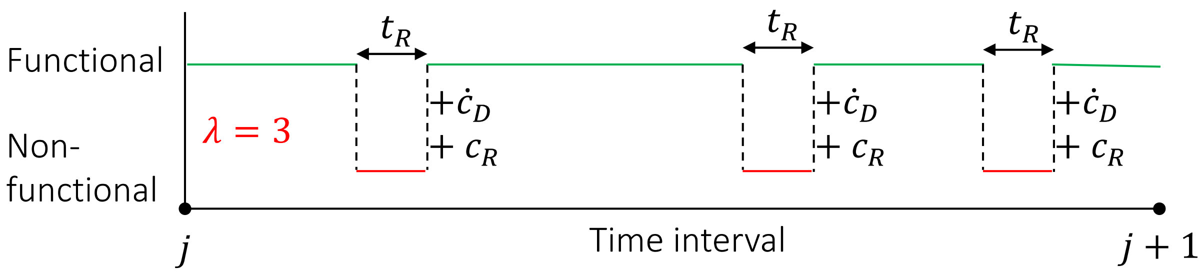

Figure 4: In a given time interval, j to j+1, each time a failure occurs, the system goes to a non-functional state for a time period of t_R. Additional cost during downtime (which is time dependent), $\dot{c}_D$, and cost to reboot are put into the system to bring it back to a functional state

Therefore, the value of the system for month j when decision is to reboot is:

$$\hat{V}_{R_j} = \dot{U}\left(T_{BH} - \hat{\lambda}_j t_R\right) - \left(\dot{c}_O\left(T_{BH} - \hat{\lambda}_j t_R\right) + \hat{\lambda}_j t_R \dot{c}_D + \hat{\lambda}_j c_R\right) \tag{8}$$

Where $\dot{U}$ is the revenue per block hour for the fleet, T_{BH} is the total block hours of the fleet for the given month, and $\dot{c}_O$ is the operating cost per block hour.

We calculate NPV when rebooting, NPV_R, at a decision point by accounting for value generated in the months after month j:

$$\widehat{NPV}_{R_j} = \sum_{i=j+1}^{N_{life}} \frac{\hat{V}_{R_i}}{(1+r)^{i-j}} \tag{9}$$

The estimated failure rate, $\hat{\lambda}_j$, is different for each month. Figure 5 shows an example of the estimated failure rate at decision point 3 (end of month 3) when the decision is to reboot. At decision point 3 we have three actual failure rates recorded for the prior months. The underlying failure rate, $\lambda_{true} = 15$, is also shown. The failure rate for the future months is estimated as the mean of the recorded failure rate. The estimated failure rate along with its uncertainty is assumed to be the same for all future months.

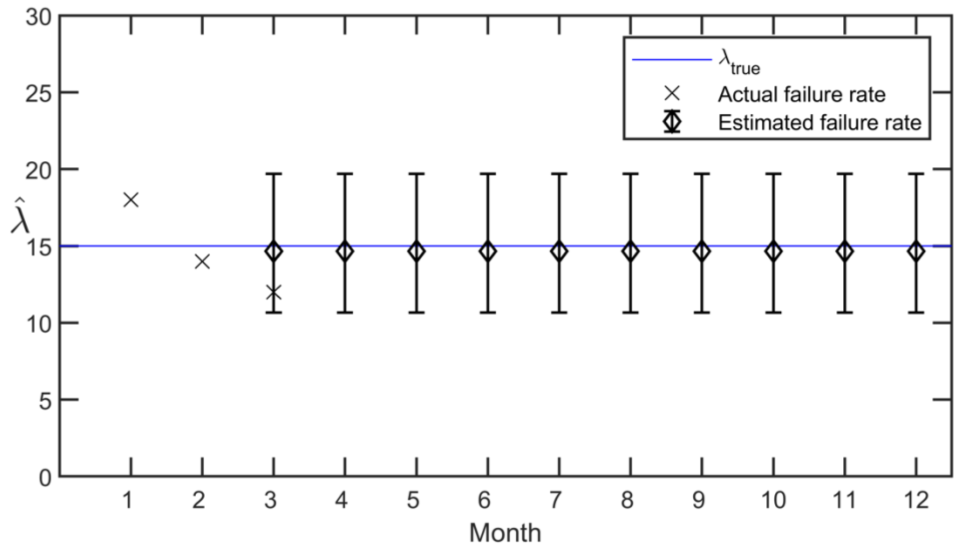

Figure 5: At the end of month 3, we assume the estimated failure rate along with its uncertainty remain same and repeat for all future months

Using the mean, upper and lower bounds of the estimated $\hat{\lambda}_j$ in Equation 7 and 8, we calculate $\widehat{NPV}_{R_3}$, which will also have its uncertainty as shown in Figure 6.

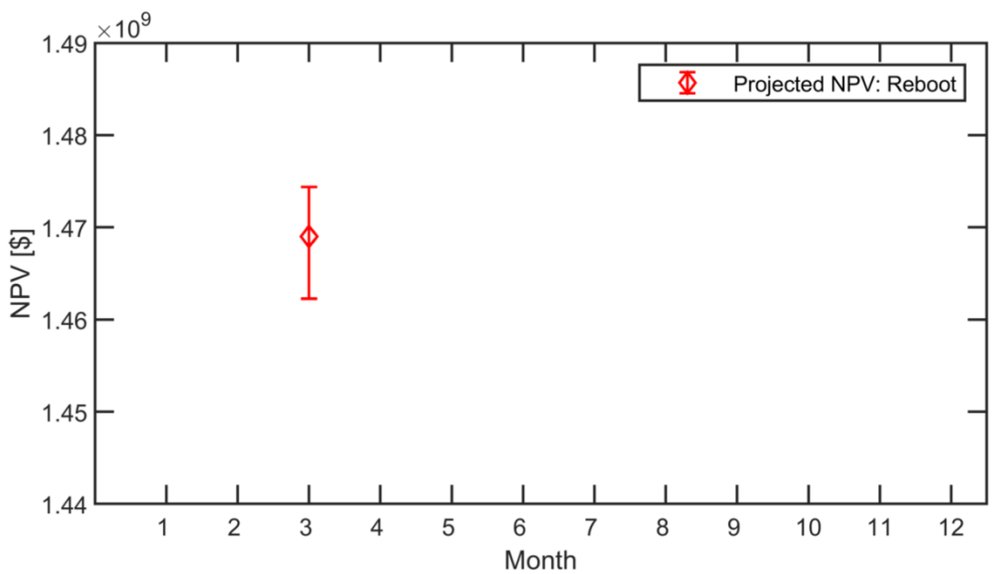

Figure 6: Using the estimated future failure rates, we calculate the NPV_R at the end of each month (month 3 shown here) and it has an uncertainty associated with it

2.2.3.2 NPV when decision is to Eliminate

By our definition, the decision to eliminate addresses the root causes of the failure mode such that the failure does not occur again. Thus, the system does not experience any recurring downtime costs. The decision to eliminate leads to a single downtime while eliminating, t_E, downtime cost per block hour, $\dot{c}_D$, and corresponding cost of elimination, c_E, which is the cost of manpower and resources required to understand the root causes and redesign the system. The downtime cost rates for both Reboot and Eliminate decisions are the same since they represent loss in revenue and additional costs while the system is not functional. The state of the system when eliminate option is chosen is shown in Figure 7.

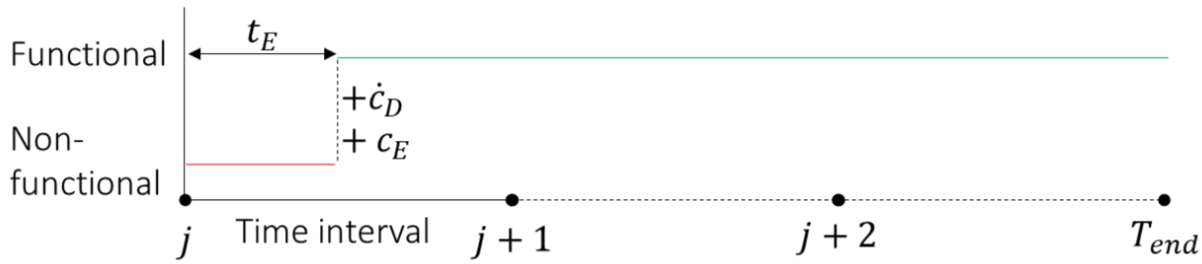

Figure 7: When eliminate option is chosen, the system experiences a downtime of t_E when the failure mode is eliminated and its associated cost, $c_E$is incurred. The downtime cost per block hour, $\dot{c}_D$

The value of the system for month j when the decision is to eliminate is:

$$\hat{V}_E = \dot{U}(T_{BH} - t_E) - (\dot{c}_O(T_{BH} - t_E) + t_E\dot{c}_D + c_E) \tag{10}$$

We calculate NPV when the eliminate option is chosen, NPV_E, at a decision point, by accounting for value generated in the months after month j. This includes the estimated value from Equation 9 and the value generated by the system in the subsequent months when it is fully functional without any recurring NFF failures:

$$\widehat{NPV}_{E_j} = \frac{\hat{V}_E}{1+r} + \sum_{i=j+2}^{T_{end}} \frac{\left(\dot{U} - \dot{c}_O\right) * T_{BH}}{(1+r)^{i-j}} \tag{11}$$

Figure 8 shows the same example chosen for the reboot option, but here the decision is to eliminate. At the end of month 3, the failure rate for the future months is zero since we assume the eliminate option removes the possibility of any failure in the chosen NFF failure mode. It follows that the failure rate is zero and has no uncertainty.

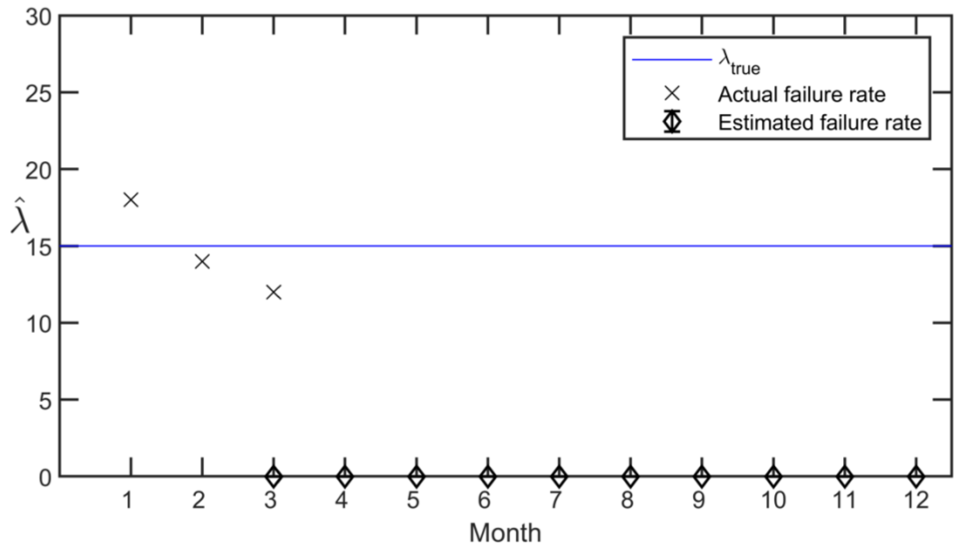

Figure 8: At the end of month 3, if decision is to eliminate, the future failure rates are zero and there is no uncertainty

Using Equation 9 and 10, we calculate $\widehat{NPV}_{E_3}$, which is shown in comparison with $\widehat{NPV}_{R_3}$ in Figure 9.

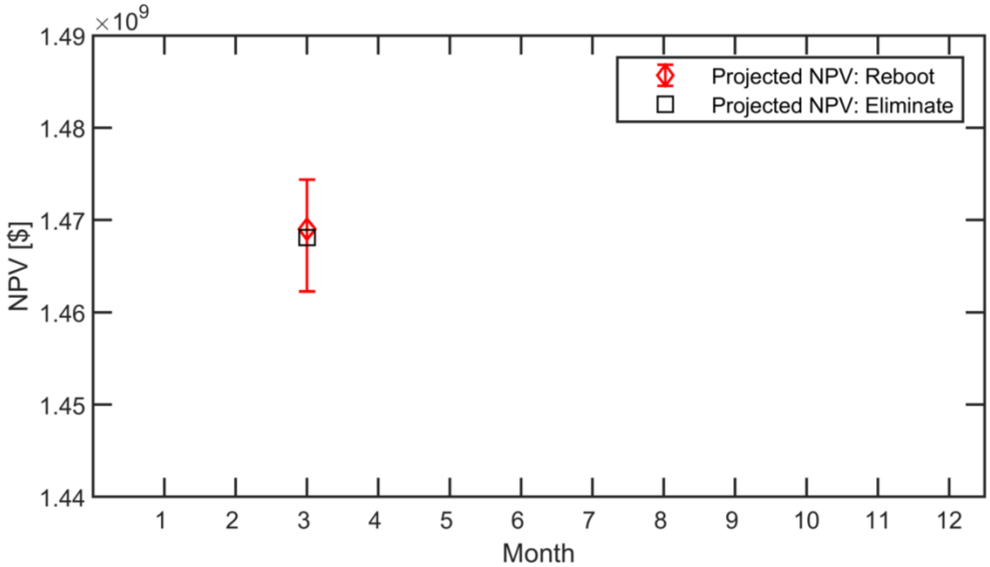

Figure 9: NPV_E is calculated assuming failure rate is zero in the future months and compared with NPV_R at the end of each month (month 3 shown here)

2.3 Uncertainty in NPV

Knowledge of the true underlying failure rate would be the ideal situation for a DM to make the correct decision at any given time. This is perfect knowledge which is unattainable but can be estimated given enough information. In the initial months, the mean of failure rate estimate might be far from the true underlying failure rate. This difference could mean that the values calculated for NPV are far from the true values of NPV. If the DM had perfect knowledge and knew the true underlying failure rate, then NPV_R will not have uncertainty bounds. In this case, NPV_E could be compared to the single certain value of NPV_R and the decision with the higher NPV could be chosen.

In the case of imperfect information, if the true failure rate is much lower than expected, then it could be better to keep rebooting till the end of aircraft system life without needing to incur the large eliminate cost. On the other hand, if the failure rate is higher than predicted, it is better to eliminate as soon as possible to reduce loss due to recurring reboot costs. Figure 10 shows two possible estimates of failure rate at the end of the first months. If in month 1, the number of failures recorded were 10, then the estimated failure rate for the future is $\hat{\lambda} = 10$ with its corresponding confidence bounds (shown in figure as Estimate 1). If the number of failures recorded were 20, then the estimated failure rate for the future is $\hat{\lambda} = 20$ with its confidence bounds (shown in figure as Estimate 2). Both these situations are possible realities due to the random nature of failures.

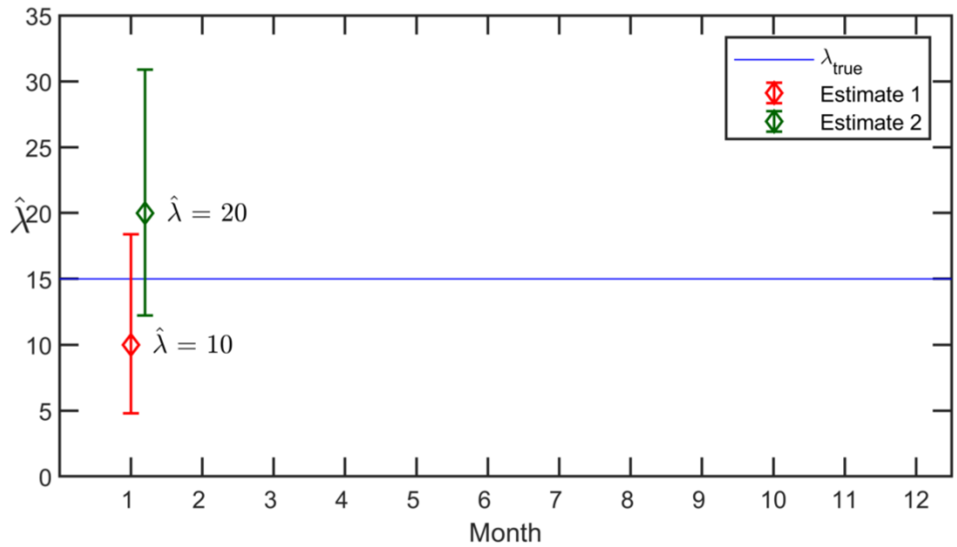

Figure 10: The estimated failure rate can be far from the underlying true failure rate and this is more likely to happen in the initial months when we have less failure rate data recorded. The two concurrent data points are separated slightly on the time-axis for readability.

The NPVs calculated for the two estimates of failure rate and λ_{true} shown in Figure 10 are plotted in Figure 11. If $\hat{\lambda}_{mean} = 10$ (estimate 1), the mean of NPV_R is greater than NPV_E, thus Reboot is the better choice if only mean values are considered. If $\hat{\lambda}_{mean} = 20$ (estimate 2), the mean of NPV_R is less than NPV_E and Eliminate is the better choice, again, if only mean values are considered. If given perfect information ($\lambda_{true} = 15$), then DM would decide to Reboot since that decision gives higher NPV.

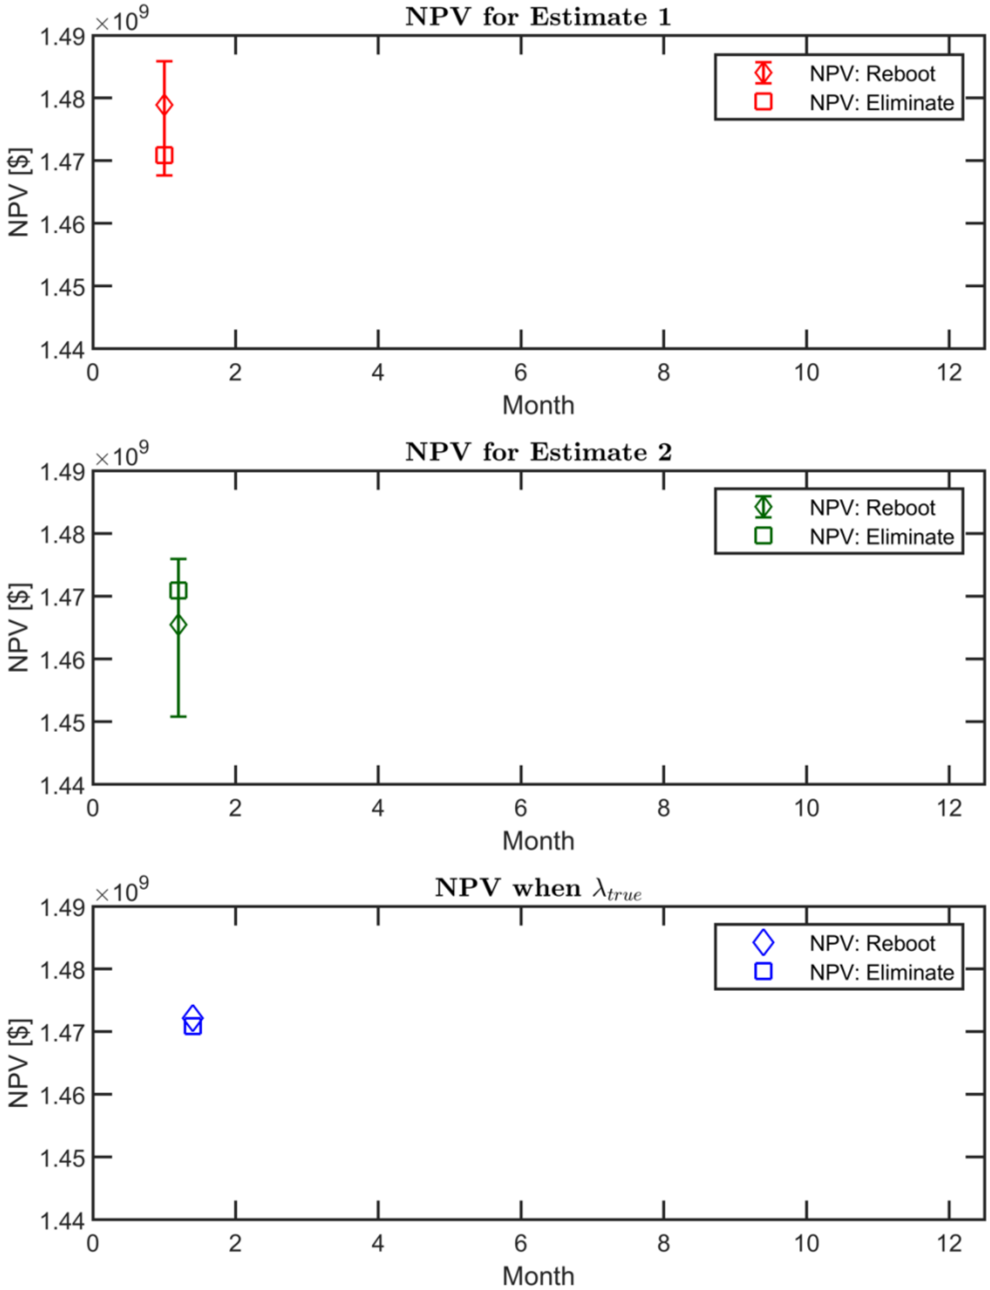

Figure 11: Comparison of NPV_R and NPV_E for the two estimates and underlying failure rate shown in Figure 14. If the DM decides without considering the uncertainty, then it could lead to a non-optimal decision

Due to NPV_E being within the uncertainty bounds of NPV_R, the DM might decide to wait and collect information to be more certain about the true failure rate but a longer wait means there is less time there is to less time to reap the benefits of not having NFFs and consequently, it is less worthwhile to eliminate. The DM needs to trade-off between the need to be certain with the need to maximize expected net present value.

Further complicating the decision is the uncertainty of the cost to eliminate and the cost per reboot. The cost to eliminate can be difficult to estimate since the NFF failure characteristics are not understood and we might not know how much it costs to fully eliminate a failure mode. The cost per reboot, as elaborated in Erkoyuncu et al. (2016) can also be difficult to estimate since there might not be enough data for such an estimate and because, within an organization, there may be multiple stakeholders in the maintenance process – the ground crew, maintenance team, engineering team, OEM, airline management – interacting with the process differently.

We incorporate these uncertainties while calculating Utility of the Eliminate decision in the next section and the results follow in Chapter 3.

2.4 Expected Utility: Quantifying DM's attitude towards uncertainty and risk

As shown in the previous section, there is a possibility of making an undesirable choice and deriving less than optimal value when the decision is based on expected NPV. A DM who is highly averse to losing money would not choose to Eliminate when uncertainties are high since there is higher risk of losing value. The decision is riskier when the uncertainties are higher. Here, we use utility theory to incorporate the DMs attitude towards risk under the uncertainties of failure rate, cost to eliminate, and cost per reboot in our decision framework.

The concept of utility is used as a measure to model the risk preference of a decision maker. Consider a 50-50 probability lottery option provided to the DM, where the win value is \$100 and lose value is -\$10 (Figure 12). The expected value of this lottery is $E(V) = 0.5 * \$100 + 0.5 * -\$10 = \$45$. A decision maker might be averse to the possibility of losing \$10. Therefore, they might accept \$25 instead of opting to bet on this lottery and win \$100. This decision attitude is termed risk-averse and the risk premium in this case is $\$45 - \$25 = \$20$. A risk-averse decision maker assigns higher utility to a profit value than what it is worth. For a risk-neutral decision maker, the risk premium is $0, which means they are neither for, nor against making a loss. A risk-seeking DM, the risk premium is negative.

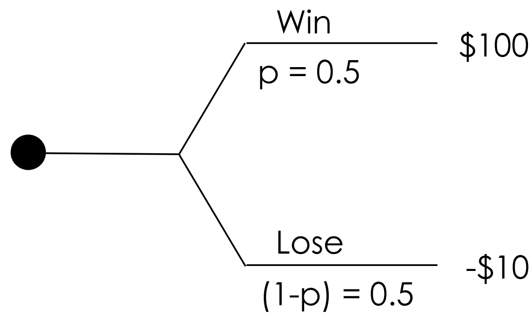

Figure 12: The expected value of this lottery is $45. The DM might accept $25 instead of choosing to take this lottery because they are averse to the possibility of losing $10

A risk-averse DM will have a decreasing rate of utility, leading to a concave utility function (Figure 13). A risk-seeking behavior is shown by a convex curve and a risk-neutral decision maker has a straight line as a risk preference function. In our work we write a utility function by using an exponential function with a single risk-aversion coefficient, γ (Buede & Miller, 2016). We define the utility function as:

$$u(x) = \frac{1 - e^{-\gamma\left(\frac{x}{x_{max}}\right)}}{1 - e^{-\gamma}} \tag{12}$$

Where γ is the risk aversion coefficient of the utility function. At x_{max}, the largest profit is expected and the utility $u(x_{max}) = 1$ and for a profit of zero, $u(0) = 0$. A more risk averse DM will have higher risk aversion coefficient, γ.

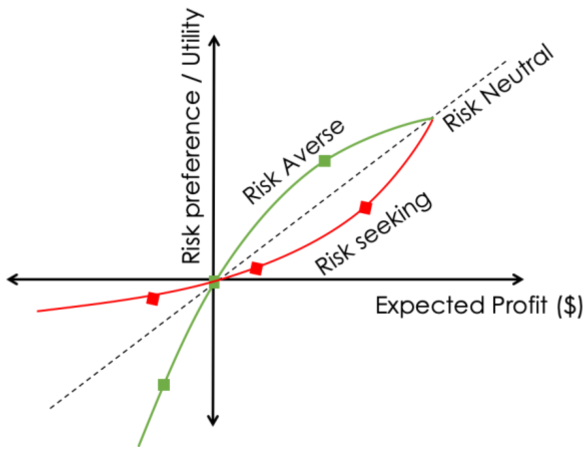

Figure 13: A risk-averse DM has a concave utility function and a risk-seeking DM has a convex utility function

In our case, the profit a DM can make with a decision to Eliminate is the difference between NPV_E and NPV_R. At each decision point, we obtain a distribution for Profit from Eliminate decision, $<NPV_E - NPV_R>$, by doing a Monte Carlo simulation with 10,000 runs over three uncertain variables, Cost to Eliminate (c_E), Cost to Reboot (c_R) and estimated failure rate ($\hat{\lambda}$) (Figure 14). We chose to do 10,000 runs after validating that the value of NPV_R converged within this number. We chose to do 10,000 Monte Carlo runs because NPV_R was seen to converge within these runs. The convergence of the simulation is shown in Appendix A.

In the absence of a good model for maintenance cost distribution, the Cost to Eliminate and Cost to Reboot are simulated to have a uniform distribution over a range from minimum to maximum possible costs. This uniform distribution is a simplifying assumption and in future can be changed to reflect more realistic cost uncertainty distributions. The minimum and maximum values used for this example are:

$$\$800 < c_R < \$3,200$$

$$\$20,000 < c_E < \$218,000$$

Failures are Poisson distributed with mean failure rate $\hat{\lambda}_{mean}$. Each Monte Carlo run calculates NPV using the random values generated from these three distributions. The profit value (Profit = $< NPV_E - NPV_R >$) calculated will also have a distribution. An example Monte Carlo simulation of 10,000 runs is shown in Figure 14. The three uncertain variables and their distributions are shown along with the resulting Profit distribution.

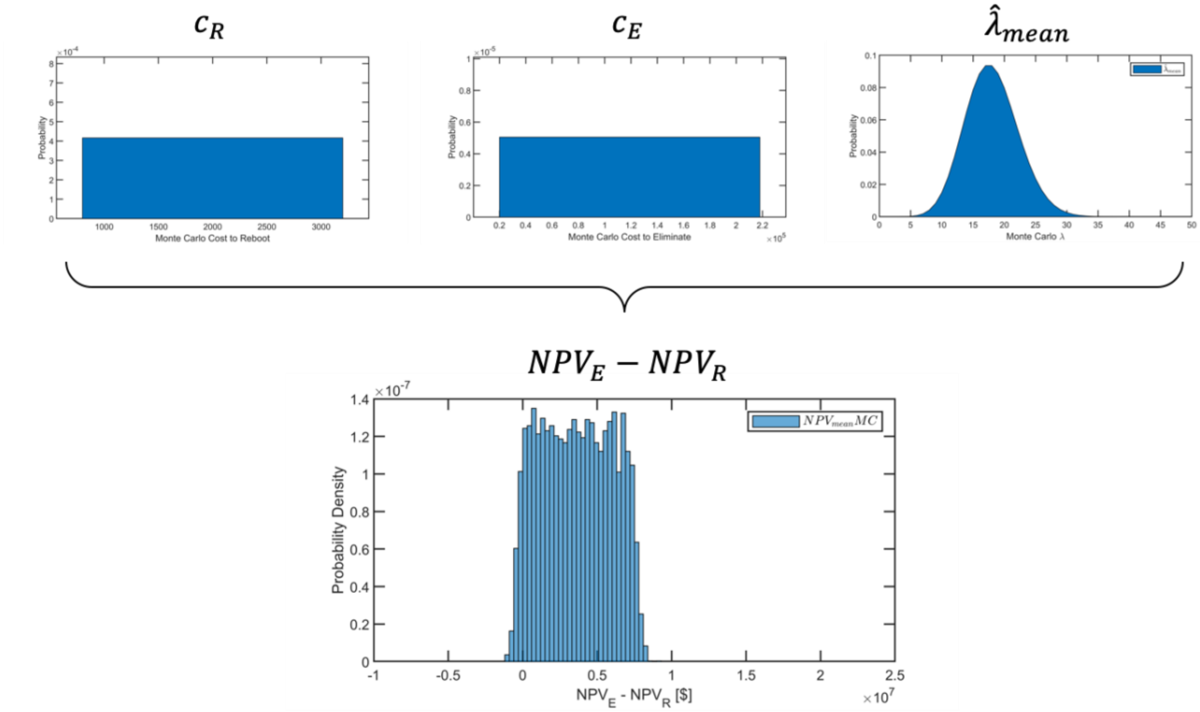

Figure 14: Cost to Reboot, c_R, and Cost to Eliminate, c_E, have a uniform distribution over their range of uncertainties. λ for each month in the future is Poisson distributed about the $\hat{\lambda}_{mean}$ of the current month. A Monte Carlo simulation over these variables gives a distribution of Profit from Eliminate decision at the current decision point

In the above example we consider only the stochastic uncertainty of failures about the estimated mean failure rate. We consider variation in $\hat{\lambda}_{mean}$ between the upper and lower bounds of the confidence interval (Equation 6 and 7). We do two more Monte Carlo simulations with 10,000 runs each using the same distributions for cost uncertainties as the previous example but with two different values for the failure rate, the upper bound of the 95% confidence interval of $\hat{\lambda}$ and the lower bound of the 95% confidence interval of $\hat{\lambda}$.

NPVs calculated from three Monte Carlo simulations gives us three different distributions of Profit from Eliminate decision (Figure 15).

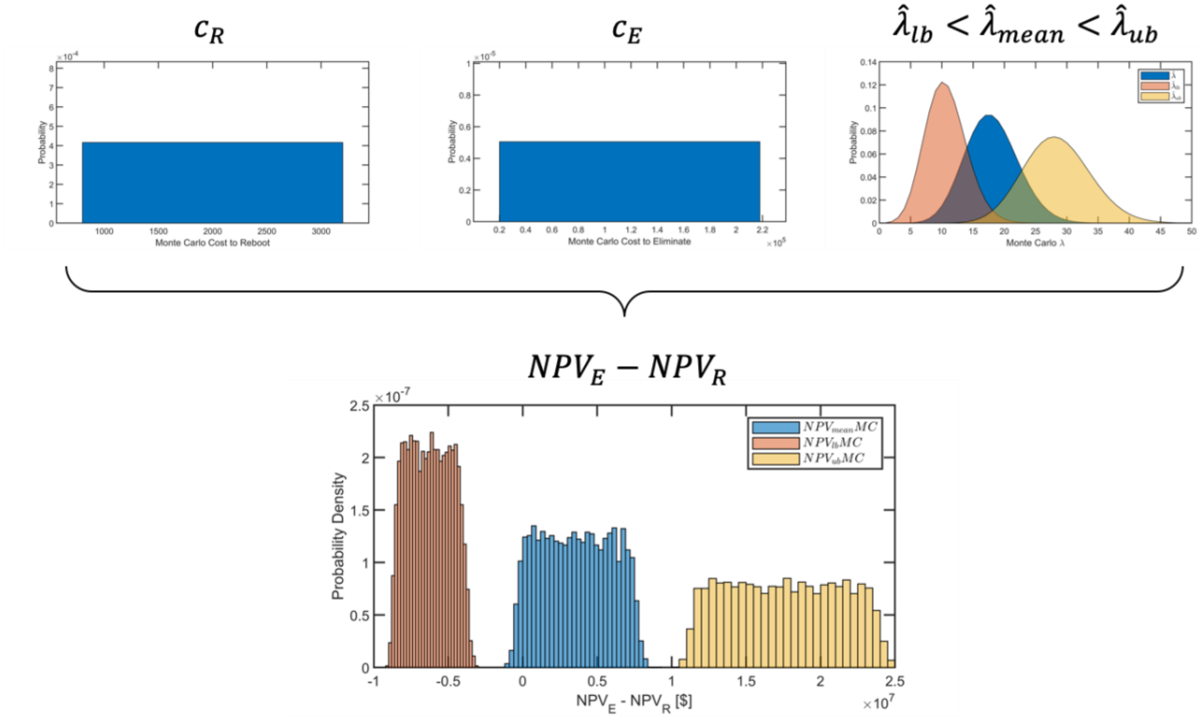

Figure 15: The Monte Carlo is repeated for the upper and lower bounds of estimated failure rate and we get three distributions of Profit from Eliminate decision

The Expected Utility, $E[u(x)]$, of choosing to Eliminate is then calculated using the probability distribution of Profit and the utility function:

$$E[u(x)] = \int_{-\infty}^{\infty} f(x)u(x)dx \tag{13}$$

Where $f(x)$ is the probability distribution of x; $x = < NPV_E - NPV_R >$ is the expected profit from the eliminate decision.

The Utility function in this example is calculated using Equation (12) with a risk-aversion coefficient $\gamma = 1$, meaning that that DM is risk-averse. Drawing an analogy to Figure 12 using the Profit values from Figure 15 where $x_{max} = \$25\, million$, the assumption of $\gamma = 1$ means accepting a situation with 0.5 probability of Profit of $x_{max} = \$25\, million$ and 0.5 probability of loss of $x_{max}/2 = \$12.5\, million$ is equal to walking away from the situation and having the certainty of neither a loss nor a profit.

The three distributions of $< NPV_E - NPV_R >$ we get from the Monte Carlo are the profit distributions, $f(x)$. Using Equation (13) we calculate three values of $E[u(x)]$ which are the mean, upper and lower bounds which the DM can use to make a decision.

2.5 Maximizing Utility in the presence of NFFs

Figure 16 shows the possible paths the system can take depending on the two decisions that a Decision Maker (DM) can take at the end of each month. At the beginning of the system lifetime, we have an unknown failure rate of faults that result in NFFs. At the end of the first month, the maintenance team will have recorded the number of NFF failures for that month. Using this data, the DM estimates a failure rate for the future months and calculates NPV_{R_1}, NPV_{E_1} and Expected Utility.

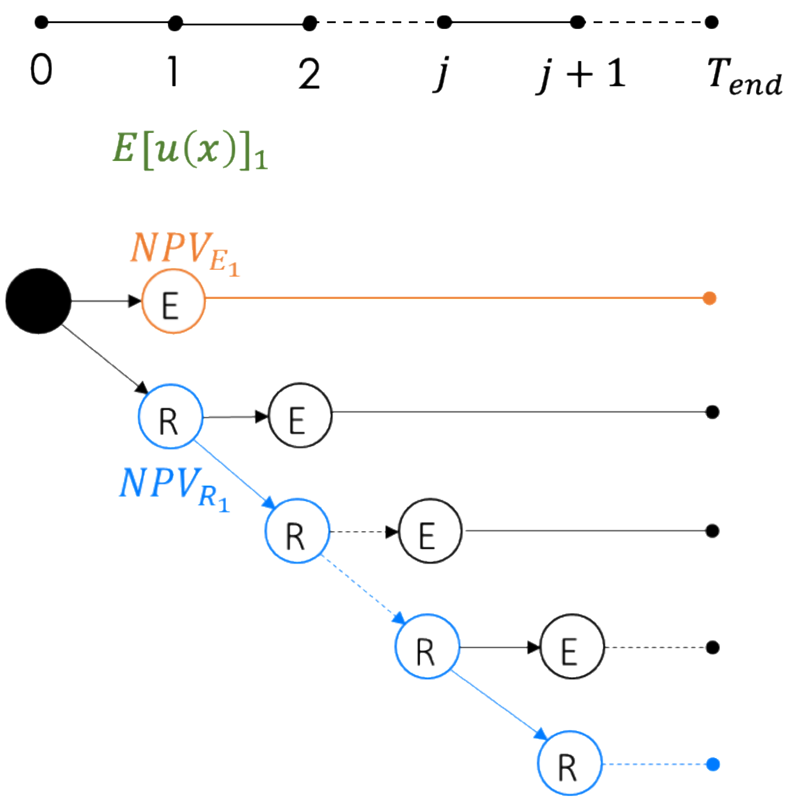

Figure 16: The system can take many paths depending on the decision made at each decision point. At the end of month 1, the DM can use NPV_E and NPV_R to determine the expected value maximizing route

If $E[u(x)] > 0$, then Eliminate is the better decision at that decision point. In Figure 17, Expected utility of the Eliminate Decision is shown for month 1. If the DM considers the lower bound of

$E[u(x)]$ for the decision, then they would decide to continue to Reboot since $E[u(x)]_{lb} < 0$. If DM considers the upper bound or mean of Utility, then, since $E[u(x)]_{mean} > 0$ and $E[u(x)]_{ub} > 0$, eliminating is the better decision and is expected to give a profit.

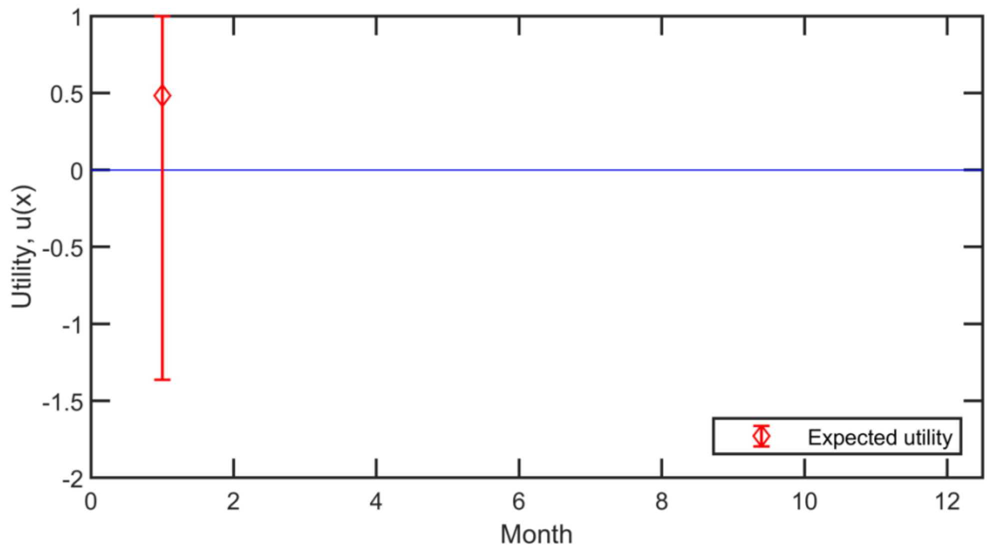

Figure 17: At the end of month 1, the DM can see the Expected Utility if they choose to Eliminate. If $E[u(x)] > 0$, then Eliminate is the better decision

At the end of the first month, if the DM chooses to continue to reboot, then the NFF failures are recorded for another month. At the end of the month 2, we have an updated failure rate estimate which is used to calculate NPV_E, NPV_R and $E[u(x)]$. Expected utility is then used to make the decision for the subsequent month (Figure 18). If the reboot option is chosen again at this point, then the same process is repeated at the end of the next month.

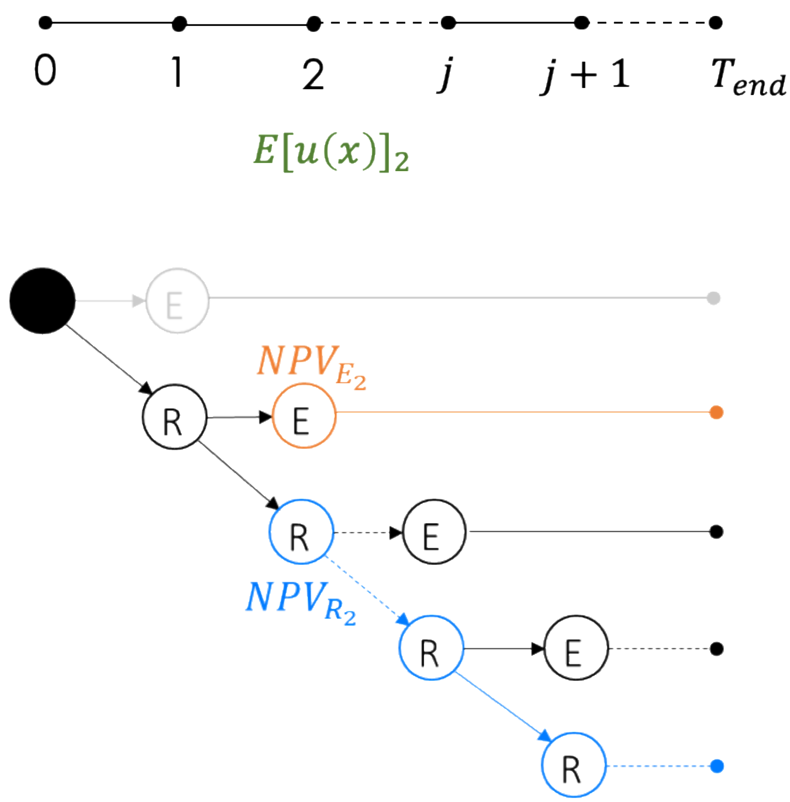

Figure 18: At the end of the first month, if the DM chooses to continue to reboot, then the NFF failures are recorded for another month. At the end of the month 2, we have an updated failure rate estimate which is used to calculate NPV_E and NPV_R and then used to make the decision for the subsequent month

2.6 Summary

We have laid out the decision framework for the two options considered – Eliminate and Reboot. An airline fleet is modeled as a revenue generating system. Failures are modeled as a stochastic Poisson process and the two decisions are assumed to have different consequences on failure rate. The NPV is calculated for both decisions but this NPV has uncertainty that can lead to undesirable decisions. There is higher risk in deciding when uncertainties are higher. The risk attitude of the DM is modeled using utility theory. The Expected Utility of Eliminate decision is calculated while accounting for uncertainties of failure rate, cost to eliminate and cost to reboot. DM can make a decision using the values of Expected Utility of Eliminate decision.

3. ANALYSIS AND RESULTS

This chapter has three parts. First, we establish a baseline result of Expected Utility and understand how it changes with time and underlying failure rate. We then explore the effect of Costs to Reboot and Eliminate on the Expected Utility and consequently on the optimal decision. We introduce a concept of threshold failure rate ($\lambda_{threshold}$) which consolidates the results such that, under uncertainties of cost and failure rate, the DM can make a decision by comparing the estimated failure rate at the current decision point with the $\lambda_{threshold}$. Finally, we show how NPV and Expected Utility changes with change in discount rate, r, and risk aversion coefficient, γ.

3.1 Input parameters for cost model

To calculate NPV and Expected Utility, the DM requires the necessary inputs for revenue, costs and failure rate. We list the required inputs in Table 1.

Table 1: Baseline numbers for input parameters required for the cost model to calculate NPV and Expected Utility

Nomenclature		Value	Unit	Source
Time interval	-	1	Months	Similar to industry practice (Kinnison & Siddiqui, 2013)
Total lifetime	T_{end}	360	Months	Follows from our assumption of an aircraft's life.
Revenue	$\dot{U}$	4,670	USD/Hour	Calculated from the operating cost using an operating margin of 6%. This percentage of operating margin is the average percentage across domestic airlines in the first quarter of 2019, as published by the Bureau of Transportation Statistics (BTS, 2019)
Operating cost	$\dot{c}_O$	4,400	USD/Hour	Taken from the value for the United airlines fleet of A320 aircraft published in Belobaba et al. (2009).
Cost during downtime	$\dot{c}_D$	10,000	USD/Hour	The Costs during Downtime are set to a reasonable value based on literature

				(Erkoyuncu et al., 2016; Hölzel et al., 2012).
Discount factor	r	5	% annually	Our approach uses a constant rate of discounting for calculating NPV, which is the common practice for engineering trade studies for systems with ~30-year life (Kenley & Armstead, 2004). The annual discount rate is converted to a monthly rate for our model.
Monthly fleet Block Hours	t_{BH}	30000	Hours	United Airlines has about 100 A320 aircraft and the average block hours per day for each aircraft is 10 hours according to the data provided by MIT's Airline Data Project (MIT, 2018). Therefore, the total monthly block hours for the A320 fleet when assuming 30 days in a month is 30,000 hrs.
Downtime per reboot	t_R	0.5	Hours	Set to a reasonable value.
Downtime to Eliminate	t_E	2,000	Hours	Set to a reasonable value.
Cost to Reboot	c_R	2,000	USD	Set to a reasonable value and is later varied to see its effect.
Cost to Eliminate	c_E	100,000	USD	Set to a reasonable value and later varied to see its effect.
Underlying true Failure rate	λ_{true}	15	per month	A failure rate of 15 NFF failures/month is a reasonable value set based on pilot reported avionics errors (Kinnison & Siddiqui, 2013). It is later varied to see its effect.
Risk aversion coefficient	γ	1	-	Set to a reasonable value and is later varied to see its effect.
Maximum Expected Profit	x_{max}	10,000,000	USD	Set to a reasonable value and is later varied to see its effect.

For our baseline result, we assume the DM is risk averse with a constant risk aversion coefficient, $\gamma = 1$ and Maximum Expected Profit, $x_{max} = \$10\ million$. We assign a utility of 1 to profits beyond this x_{max} (Figure 19).

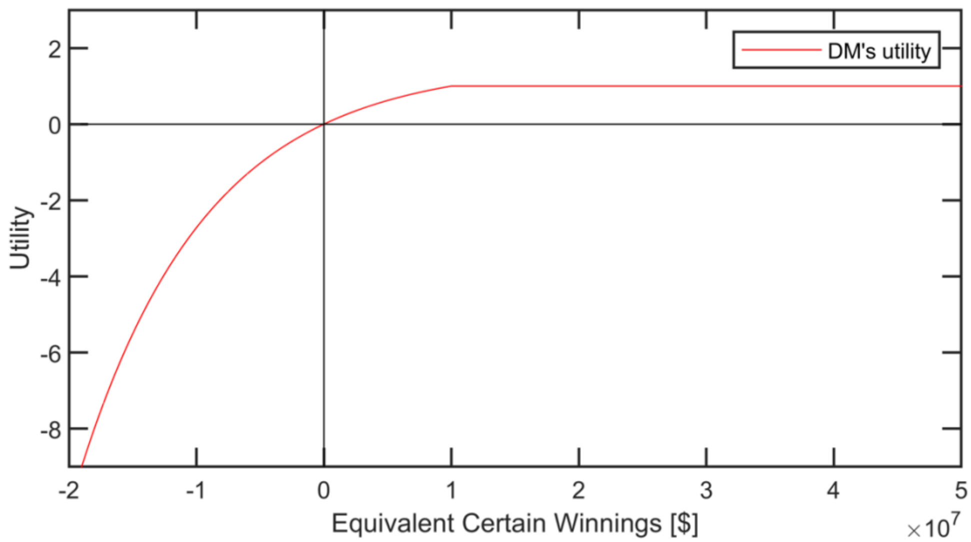

Figure 19: We establish baseline results with a utility function having a constant risk aversion coefficient, $\gamma = 1$ and a utility of 1 is assigned to profits beyond $x_{max} = \$10M$

3.2 Change of NPV and Expected Utility with time

We use a sequence of random failure rates simulated from a Poisson distribution about an underlying true failure rate of $\lambda_{true} = 15\ failures/month$. Figure 20 shows the failure rate sequence along with the estimated failure rate for each decision point in the first year of the analysis.

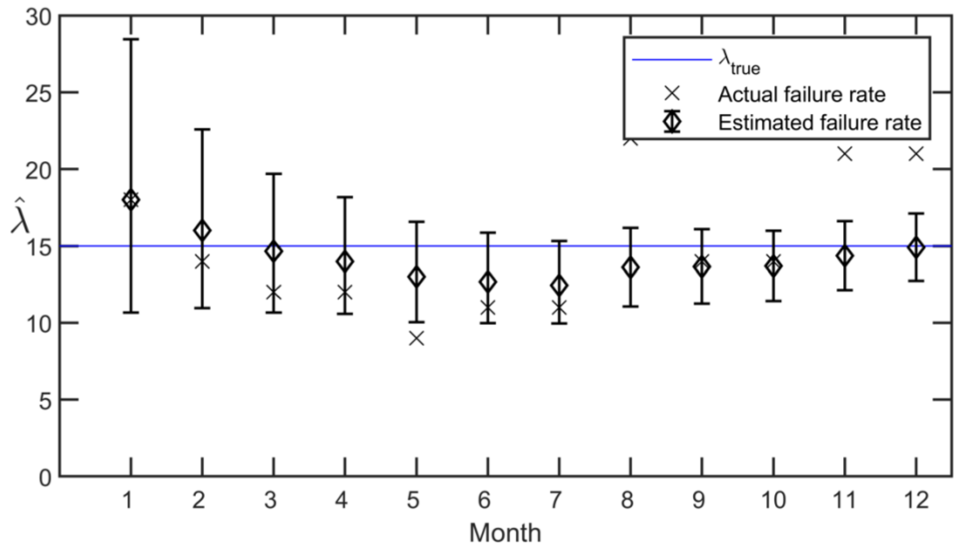

Figure 20: The following NPV and Expected Utility results use the failure sequence shown in this figure

Figure 21 shows NPV calculated for the first year using this failure rate sequence. Figure 22 shows NPV calculated for the full lifetime of the aircraft fleet. We make the following observations for the NPV calculated:

1. Over the course of the fleet's life, NPV_R and NPV_E both gradually reduce since they are calculated looking forward from that point, so there is less remaining life to reap value. $NPVs$ do not have a constant rate of decrease because of discount rate accounting for time value of money.
2. NPV_R changes with change in $\hat{\lambda}$ and therefore increases or decreases from one decision point to the next. NPV_E does not depend on λ and therefore gradually and monotonically decreases with each decision point.
3. The confidence interval of $\hat{\lambda}$ decreases with time and consequently, the confidence interval of NPV_R also becomes smaller with time.

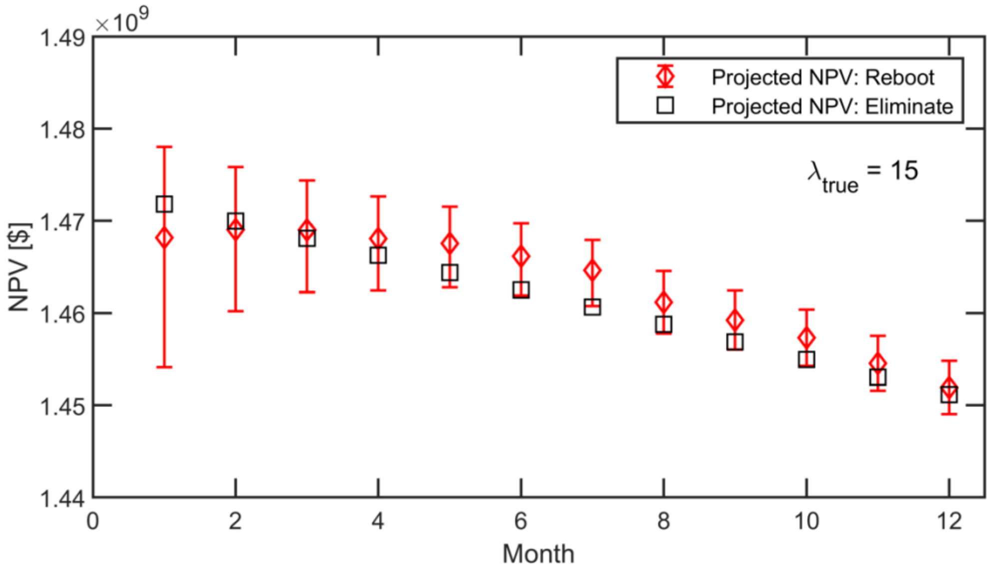

Figure 21: NPV_R and NPV_E are calculated using the failure sequence shown earlier. NPV_E has a more gradual decrease since it does not depend on λ

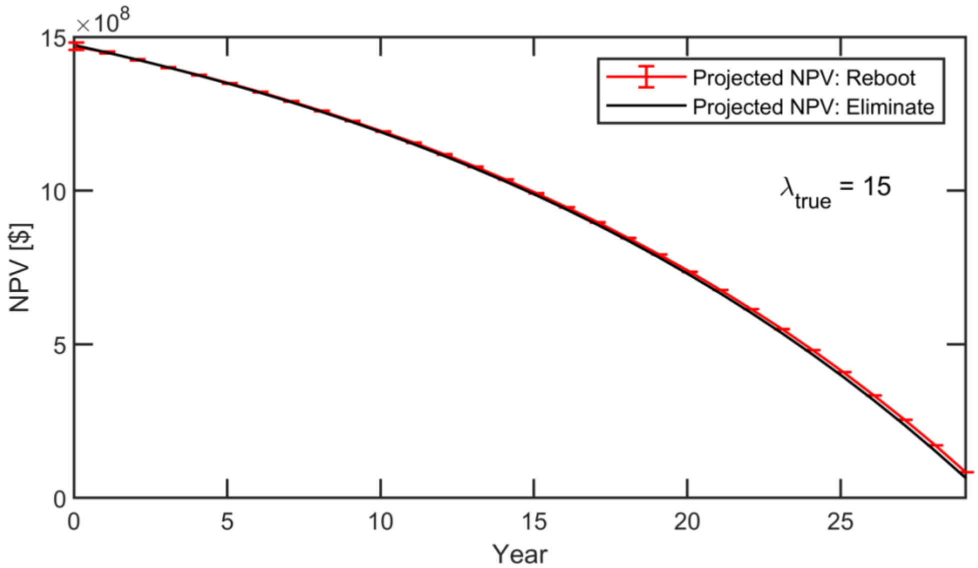

Figure 22: Over the full lifetime of the fleet, NPV reduces as the remaining life from which we derive value decreases. The values of $NPV_R$and NPV_E decrease by a large amount from beginning to end of life. The difference between the two NPVs appear small compared to the change across the lifetime

Figure 23 shows the Expected Utility for all the decision points in the first year. If the DM chooses to make a decision based on the lower bound of $E[u(x)]$ then, Eliminate is not the optimal decision since $E[u(x)]_{lb}$ is not positive for any of the decision points in the first year. If the DM chooses to decide based on the mean value or upper bound of $E[u(x)]$ then, the Eliminate decision is optimal in the first month.

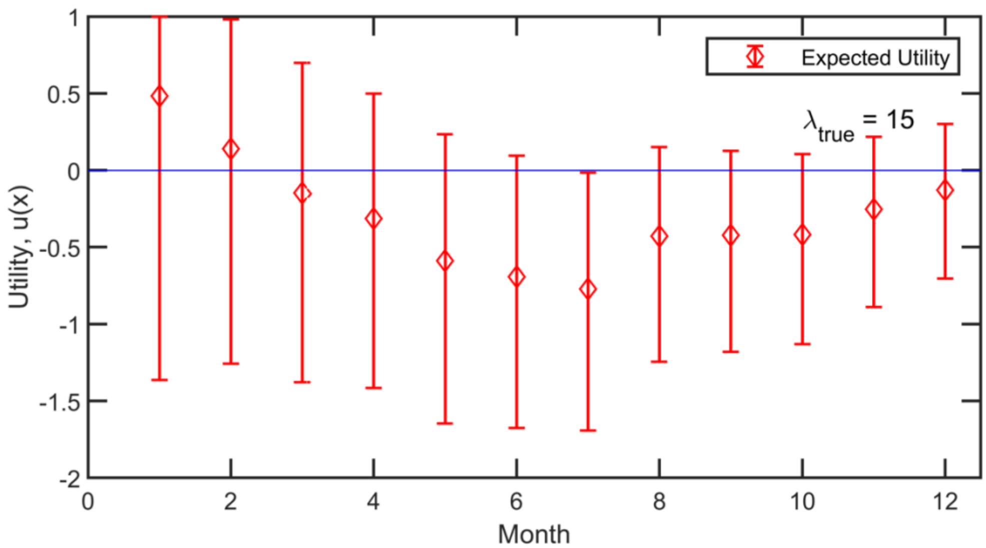

Figure 23: In the first year, the lower bound of $E[u(x)]$ does not become positive, therefore the Eliminate is not the optimal decision in the first year

Expected Utility also varies with varying $\hat{\lambda}$. If $\hat{\lambda}$ is higher, then we incur costs for more reboots per month and therefore NPV_R is lower and Profit from Eliminate decision is more. If Profit is higher, the utility in deciding to Eliminate is higher and Expected Utility, $E[u(x)]$ is more. Similar to NPV_R, the confidence interval of $E[u(x)]$ decreases with time.

As the remaining life of the fleet reduces, Expected Utility gradually decreases (Figure 24). It is less worthwhile to Eliminate as time progresses as there is less time to reap the benefits of eliminating the failure mode or in other words, we can continue to reboot and still get more value than spending the large amount of eliminate cost.

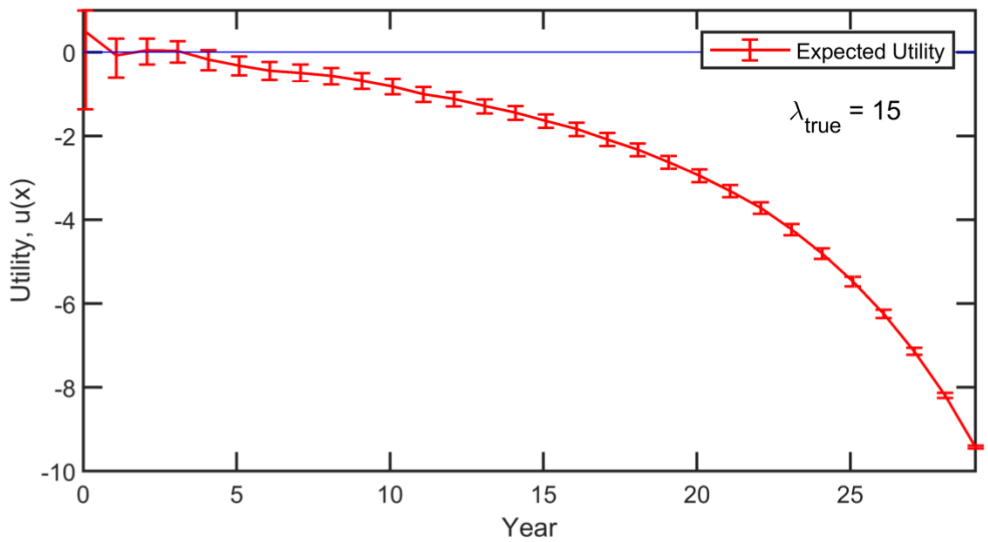

Figure 24: Expected Utility reduces as the remaining life to recover the Cost to Eliminate reduces

3.3 Change of NPV and Expected Utility with λ_{true}

If the underlying true failure rate, λ_{true}, is increased to 20 (Figure 25 shows the simulated failure rate sequence), then we can make the following observations:

1. NPV_E remains the same as baseline values since it does not depend on λ.
2. Overall, NPV_R is lower since more Reboots are needed each month (Figure 26). Consequently, profit from Eliminate decision is more and $E[u(x)]$ is higher (Figure 27).
3. The confidence interval of $E[u(x)]$ decreases with time. This increase in certainty, combined with the higher failure rate, makes $E[u(x)]_{lb} > 0$ after 12 months of collecting failure rate data and allows the DM to choose to Eliminate at that time. If the DM waits for a long time to collect more data, we have tighter confidence intervals but the $E[u(x)]$ becomes negative again as Eliminate becomes less worthwhile.

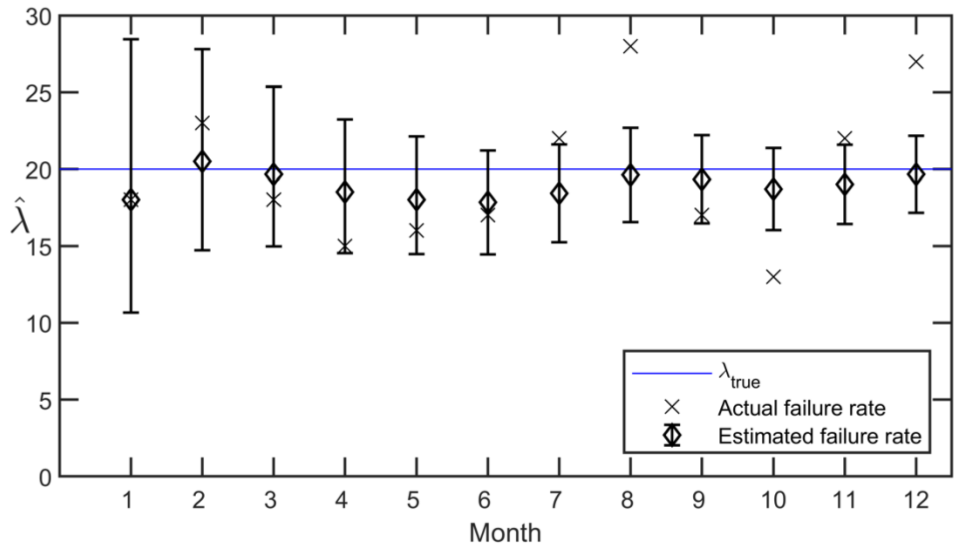

Figure 25: Simulated failure sequence for $\lambda_{true} = 20$

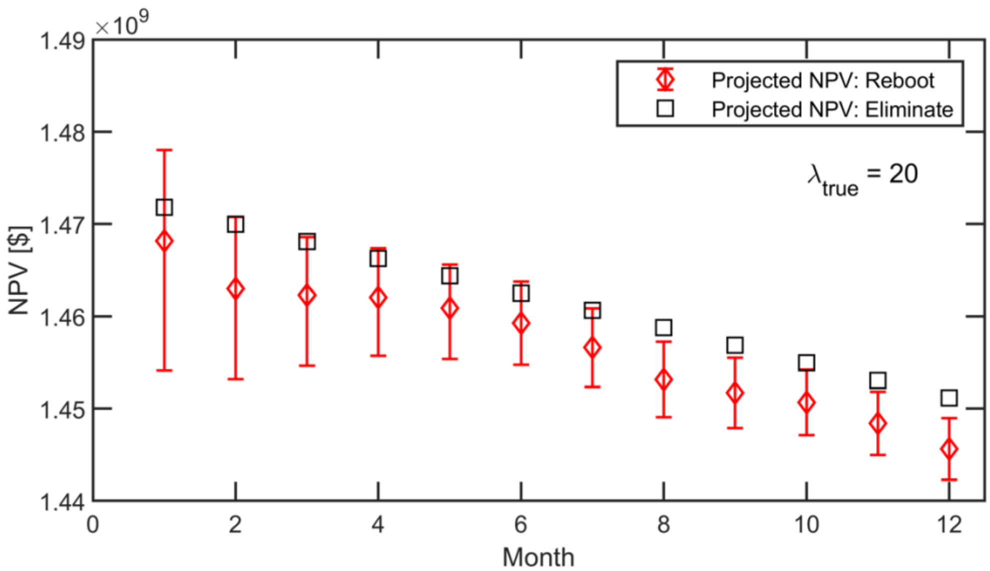

Figure 26: When comparing the results for $\lambda_{true} = 20$ with $\lambda_{true} = 15$, we find that for higher λ_{true}, NPV_R is lower

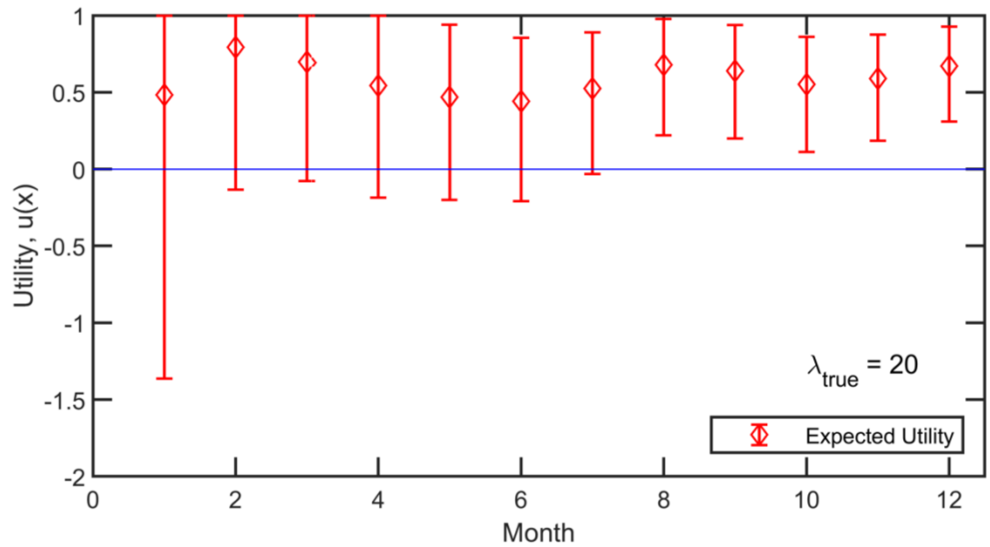

Figure 27: Expected Utility increase with increase in λ_{true}

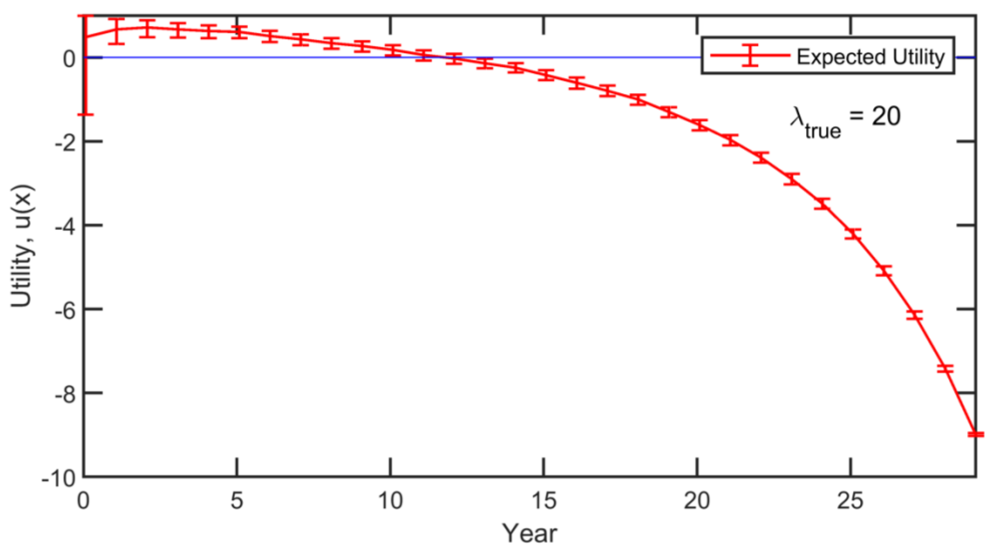

Figure 28: The Eliminate decision have more utility in the beginning of the lifetime because there is more time to recover the Cost to Eliminate. Later in the system lifetime, although we are more certain of the failure rate estimate, Eliminate is not the optimal decision

3.4 Change of NPV Discount rate, *r*

Discount rate accounts for the time value of money. A high discount rate means that investments made in the past or present grow more rapidly; conversely, payments in the future are worth less in the present than if the discount rate were lower. To illustrate this, we simulate NPV for different discount rates and compare them. For a lower discount rate ($r = 2\%$), we get high NPV in the present and for higher discount rate, we get lower NPV in the present (Figure 29). The differences between the NPVs calculated for different discount rates are higher initially and become smaller with time, because the effect of discounting increases the further into the future we look.

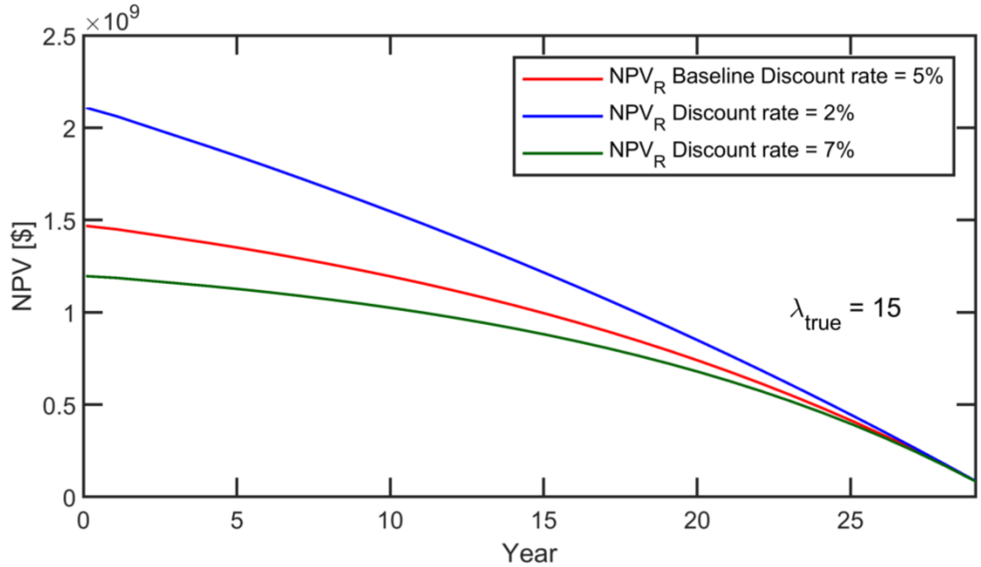

Figure 29: Higher discount rate means the value in the future are worth lesser in the present. Therefore, NPV is higher if discount rate is lower

3.5 Change of Expected Utility with risk preference

We change the Risk Aversion coefficient to model the effect of the DM's risk attitude on decision making. Figure 30 shows the utility functions for a More Risk Averse DM ($\gamma = 1$), Less Risk Averse DM ($\gamma = 0.2$) and a Risk Neutral DM (Utility function is a straight line). The Utility function mostly changes in the negative axis where a Loss is expected instead of Profit. A highly risk averse DM has very low utility for value lost (profit is negative).

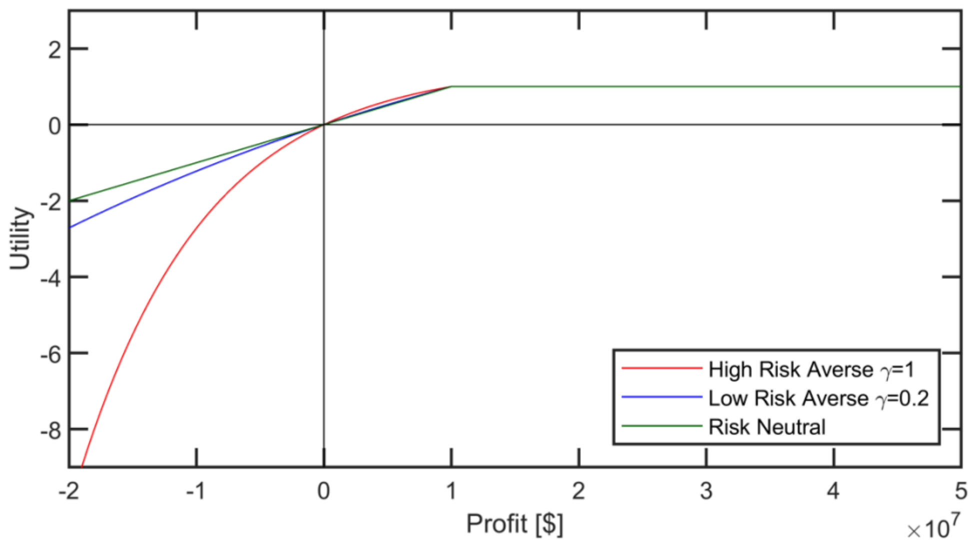

Figure 30: We change utility function to see the effect of DM's risk attitude on decision making

Figure 31 shows the Expected Utility calculated using the three different utility functions shown in Figure 30. If the Expected Profit is negative, then a Risk Averse DM will see less Expected Utility than a Risk Neutral DM. Conversely, if Expected Profit is positive, then a Risk Averse DM will see more Expected Utility than a Risk Neutral DM.

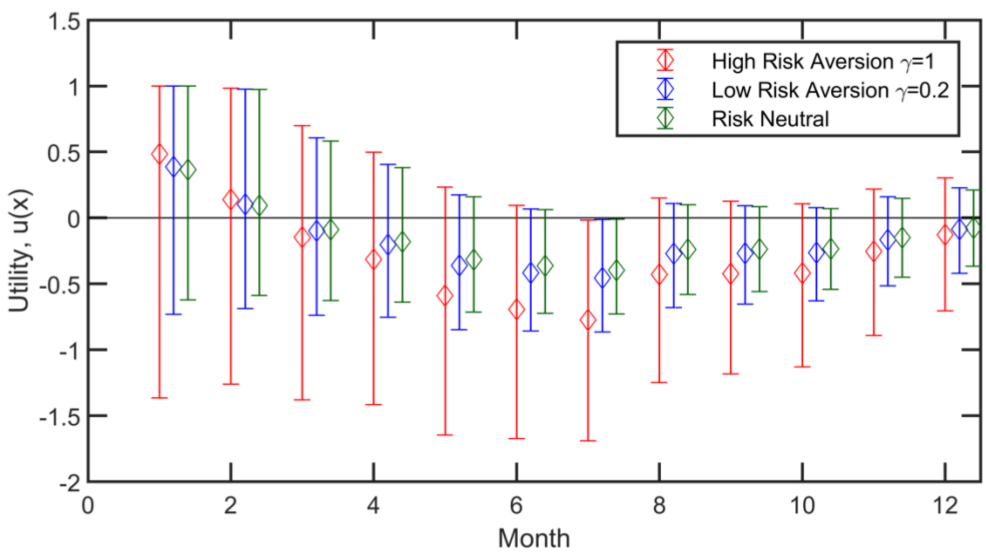

Figure 31: The lower bound of Expected Utility for highly risk averse DM is very low because DM strongly does not want loss in value

3.6 Change of Expected Utility with change in Maximum Expected Profit (x_{max})

In the Baseline simulations, we assumed that the DM expects a maximum profit of \$10 million, above which the utility is equal to 1. We change the Maximum Expected Profit to see the effect it has on Expected Utility. Figure 32 shows the utility functions for the baseline $x_{max} = \$10\ million$ and for the new $x_{max} = \$25\ million$. Both the utility functions shown are for the same risk aversion coefficient $\gamma = 1$. The Utility increases if the DM expects a loss and reduces if there is a profit. Above $x_{max} = \$25\ million$, the utility is 1 for both functions.

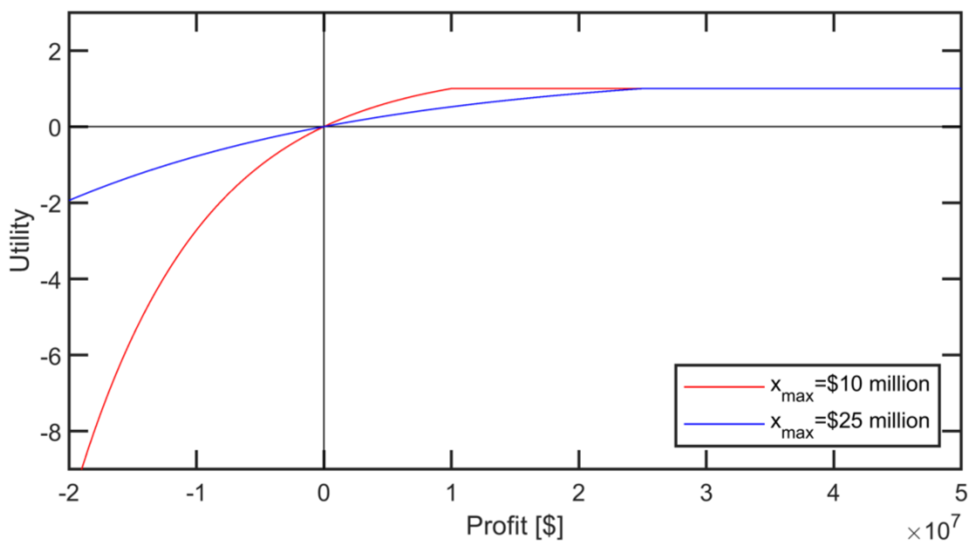

Figure 32: We change the maximum expected profit, x_{max}, and see its effect on Expected Utility calculated

Figure 33 shows the Expected Utility calculated using the two different Maximum Expected Profit values. If a loss is expected, then a DM who has a higher Maximum Expected Profit, will see it having higher utility. Conversely, if a profit is expected, then a DM who has a higher Maximum Expected Profit, will see it having lower utility.

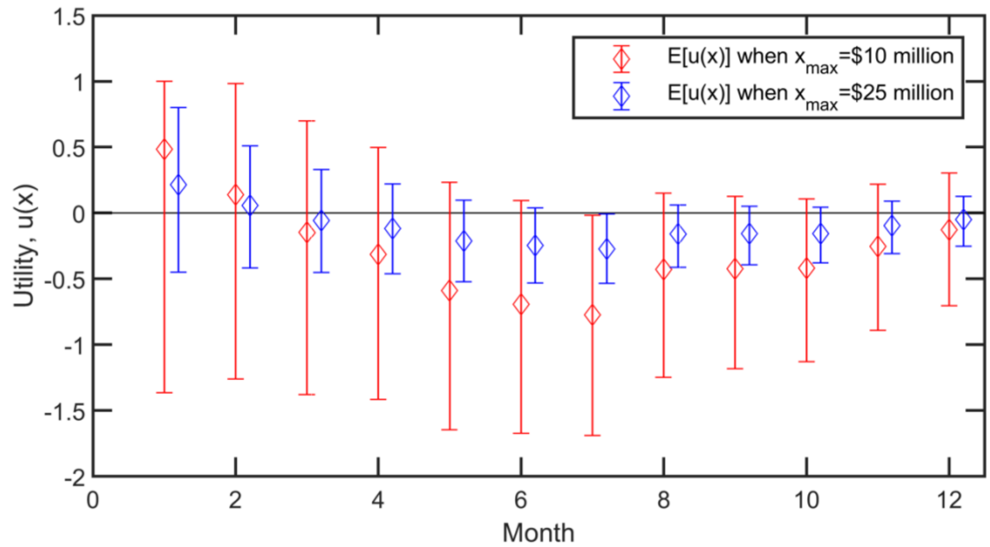

Figure 33: The lower bound of Expected Utility for DM who has higher Maximum expected Profit, x_{max}, has lower Expected Utility when making a profit and higher Expected Utility when making a loss

3.7 Change of NPV and Expected Utility with Cost to Eliminate and Cost to Reboot

If the Cost to Eliminate and Cost to Reboot are double the baseline values, ($c_R = \$4{,}000 \ and \ c_E = \$200{,}000$), then we can make the following observations:

1. NPV_E and NPV_R are lower because c_E and c_R are higher (Figure 34).
2. NPV_R is more sensitive to change in c_R compared to sensitivity of NPV_E to change in c_E. This difference in sensitivity is because the Cost to Reboot is incurred each time a failure occurs for all the months in the remaining life. The Cost to Eliminate is incurred just once when the DM decides to Eliminate.
3. $E[u(x)]$ is higher since NPV_R is lower and Profit from Eliminate decision is higher (Figure 35).

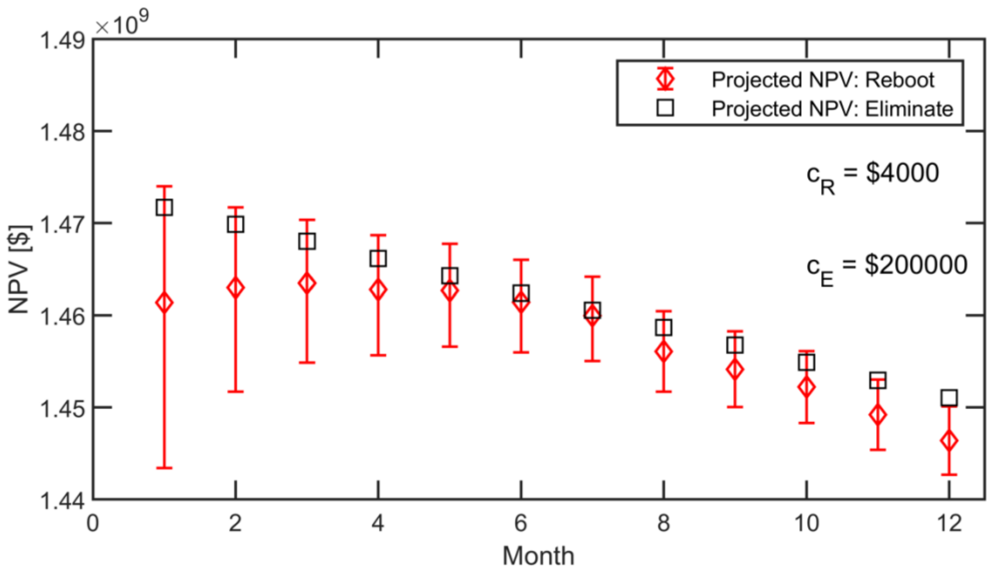

Figure 34: NPV_R is more sensitive to change in c_R than NPV_E is sensitive to change in c_E

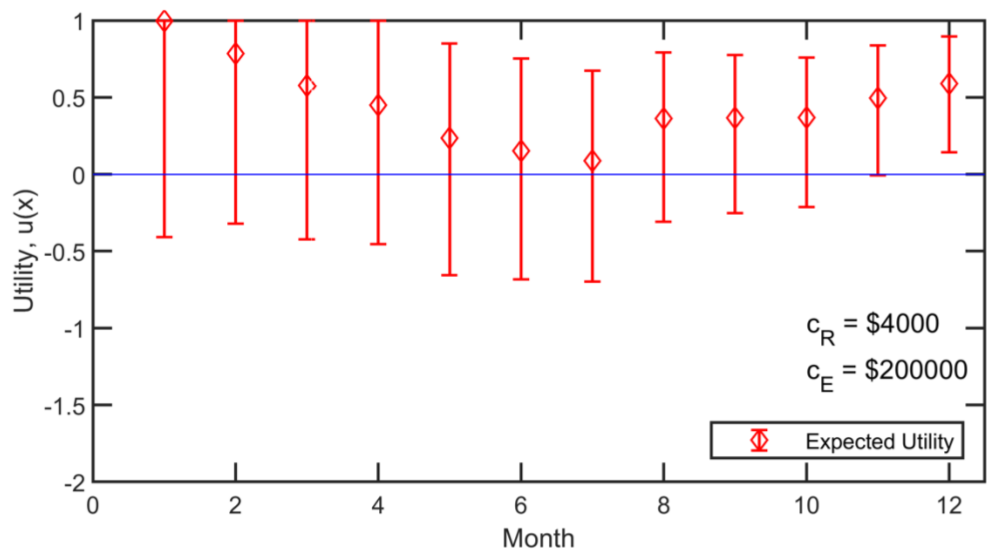

Figure 35: Lower NPV_R means there is more Profit gained from choosing to Eliminate and therefore more Expected Utility

3.8 Expected Utility for a range of Cost to Eliminate and Cost to Reboot

We extend the results in the previous section to a range of c_R and c_E costs. Figure 36 is a heat map of $E[u(x)]_{lb}$ for all c_R and c_E costs at six decision points (month 1, 5, 10, 15 and 24). More negative values of $E[u(x)]_{lb}$ are more red in the heat map, and more positive values are blue.

We make the following observations:

1. No costs in the ranges shown yield $E[u(x)]_{lb} > 0$, therefore Eliminate is not the optimal decision for any of these costs. This is because, for the cost ranges simulated, the failure rate of $\lambda_{true} = 15$ is low enough that the Eliminate decision does not yield the expected utility. Eliminate is not a worthwhile choice even after waiting to reduce uncertainty till month 24.
2. When Cost to eliminate is low and Cost to Reboot is high, the Eliminate decision has higher Profit and therefore higher Expected Utility. Conversely, for higher Cost to Eliminate and low Cost to Reboot, Expected Utility is lower.
3. In the beginning of the lifetime, $E[u(x)]_{lb}$ may decrease (month 1 to month 5) or increase (month 5 to month 24) as the estimate of failure rate is also non-monotonic with time. The same behavior is shown in Figure 23 and Figure 24 .

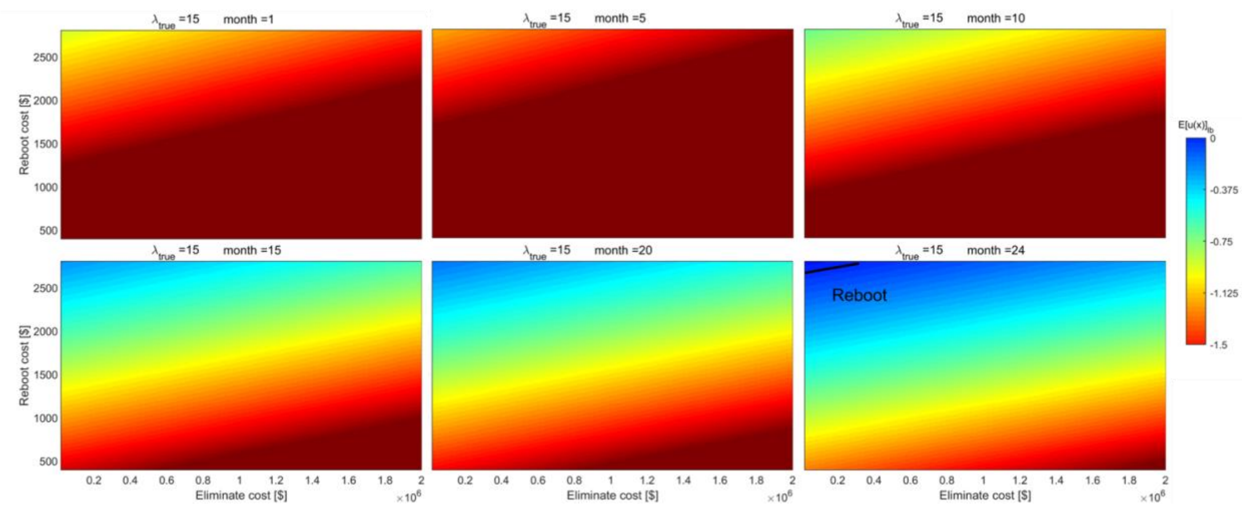

Figure 36: Heat map of $E[u(x)]_{lb}$ for $\lambda_{true} = 15$ shows that it may decrease (month 1 to month 5) or increase (month 5 to month 24) in the beginning of the lifetime. Reboot yields greater expected NPV in all cases.

If λ_{true} is higher, then Profit from Eliminate decision would be higher and Expected Utility is also higher. This is shown in Figure 37, where we plot $E[u(x)]_{lb}$ for $\lambda_{true} = 20$. The black line (where $E[u(x)]_{lb} = 0$) distinguishes the heat map into regions where Eliminate is the better decision (where $E[u(x)]_{lb} > 0$) or Reboot is the better decision (where $E[u(x)]_{lb} < 0$).

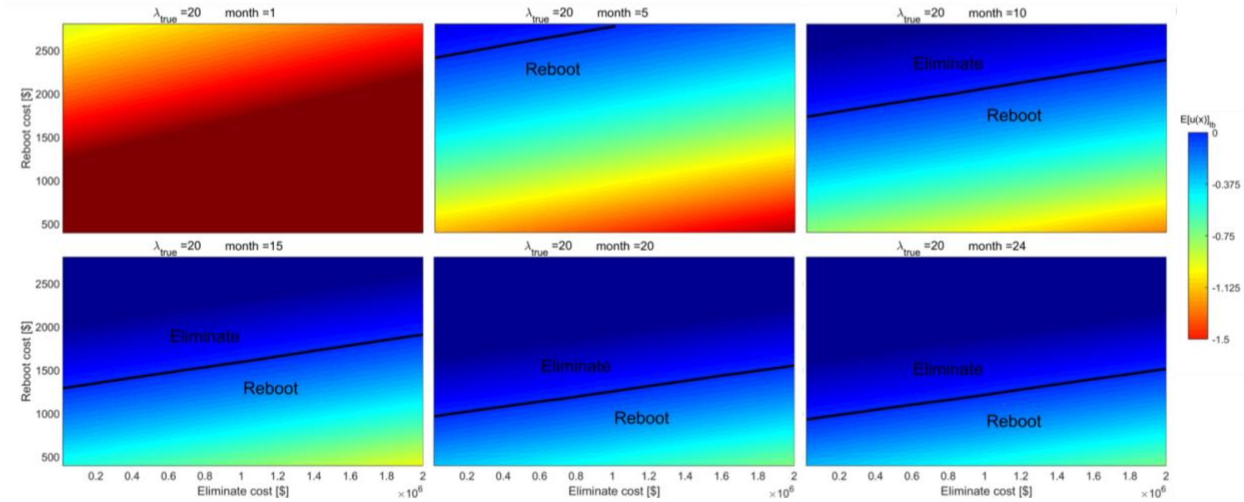

Figure 37: $E[u(x)]_{lb}$ for $\lambda_{true} = 20$ is higher than $\lambda_{true} = 15$ for all cost values in the ranges considered. The black line is the where $E[u(x)] = 0$. Reboot is preferred in month 1 and 5.

$E[u(x)]_{mean}$ (shown in Figure 38 for $\lambda_{true} = 15$)is greater than $E[u(x)]_{lb}$ and has more costs where it is greater than zero and Eliminate is the optimal choice. Eliminating is optimal even if the cost to eliminate are higher and cost to reboot are lower. If the DM chooses to use $E[u(x)]_{mean}$ to make the optimal choice, then they may be able to make the decision earlier in the lifetime since they are willing to accept more risk by not accounting for uncertainty in the estimate of mean failure rate.

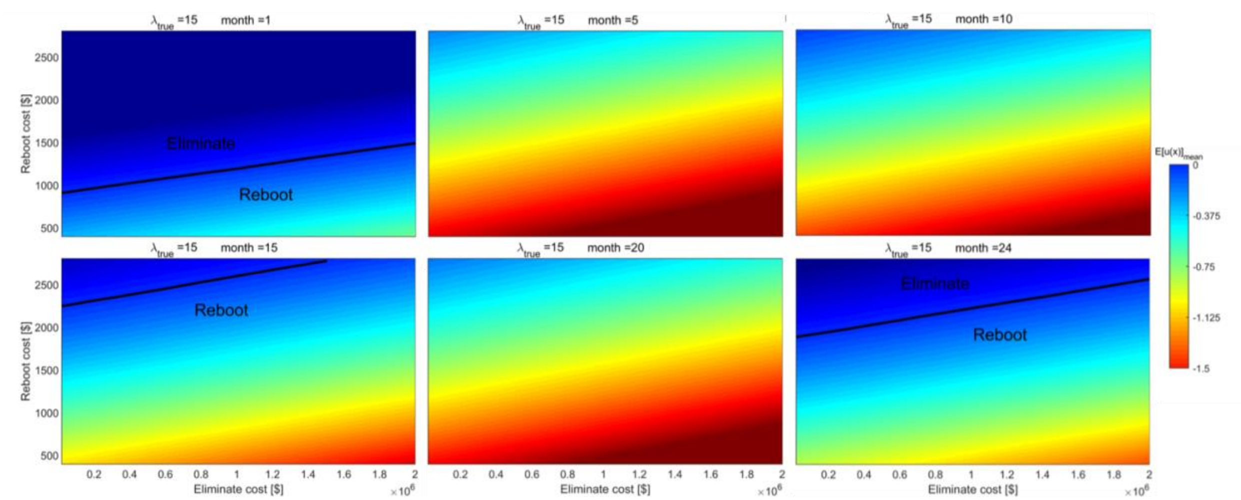

Figure 38: If the DM chooses to use the mean of Expected Utility instead of its lower bound (shown in Figure 33), then they can make the decision earlier in the lifetime since uncertainty in estimated mean of failure rate is not considered

3.9 Threshold Failure rate ($\hat{\lambda}_{threshold}$)

We introduce the concept of threshold failure rate to provide the DM with a single parameter metric to make the optimal decision. At a decision point i:

$$if\ \hat{\lambda}_i > \hat{\lambda}_{i,threshold}, Eliminate\ is\ the\ optimal\ decision$$

$$and\ if\ \hat{\lambda}_i < \hat{\lambda}_{i,threshold}, Reboot\ is\ the\ optimal\ decision$$

With the baseline cost values, we find $\hat{\lambda}_{threshold}$ for all decision points i. The mean, upper bound and lower bound of $\hat{\lambda}_{threshold}$, are values of $\hat{\lambda}$ when $E[u(x)]_{mean} = 0$, $E[u(x)]_{lb} = 0$ and $E[u(x)]_{ub} = 0$ respectively. With greater uncertainty in the beginning of the lifetime, $\hat{\lambda}_{threshold}$ also has larger bounds (Figure 39). This means if the DM uses $E[u(x)]_{lb}$, then $\hat{\lambda}_i$ should be high (>25 at Month 1 for the baseline simulation) for Eliminate to become the optimal choice. With time, even with the lower $\hat{\lambda}_i$, we get required utility from the Eliminate decision, therefore $\hat{\lambda}_{threshold}$ decreases.

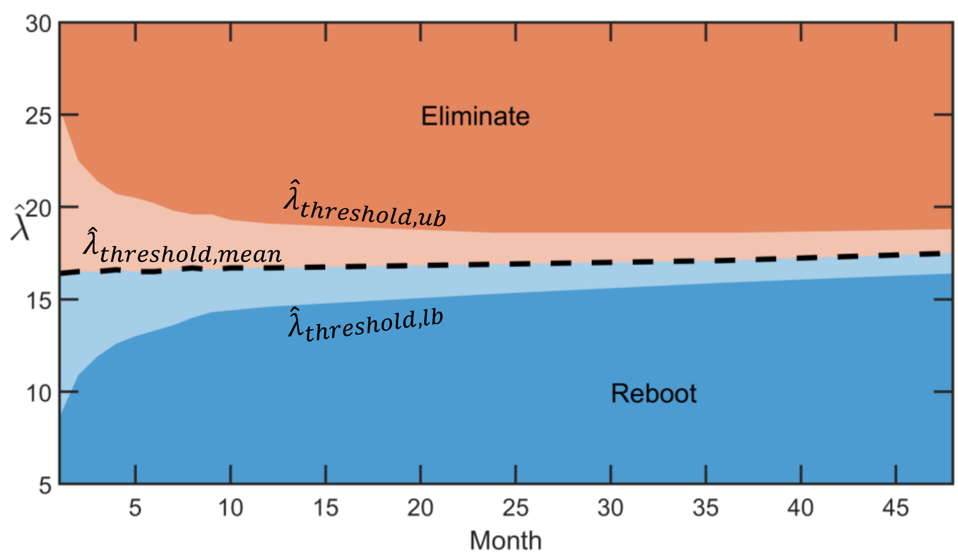

Figure 39: In the beginning of life, more uncertainty in estimated failure rate leads to larger bounds on $\hat{\lambda}_{threshold}$

The uncertainty decreases as we move further into the future but there is less time to recover the Cost to Eliminate and rebooting becomes more worthwhile. Close to the end of the fleet's life, it is better to continue to Reboot unless the failure rate is very high leading to Reboot costs becoming greater than Eliminate cost. Therefore, $\hat{\lambda}_{threshold}$ increases exponentially as we approach the end of life (Figure 40).

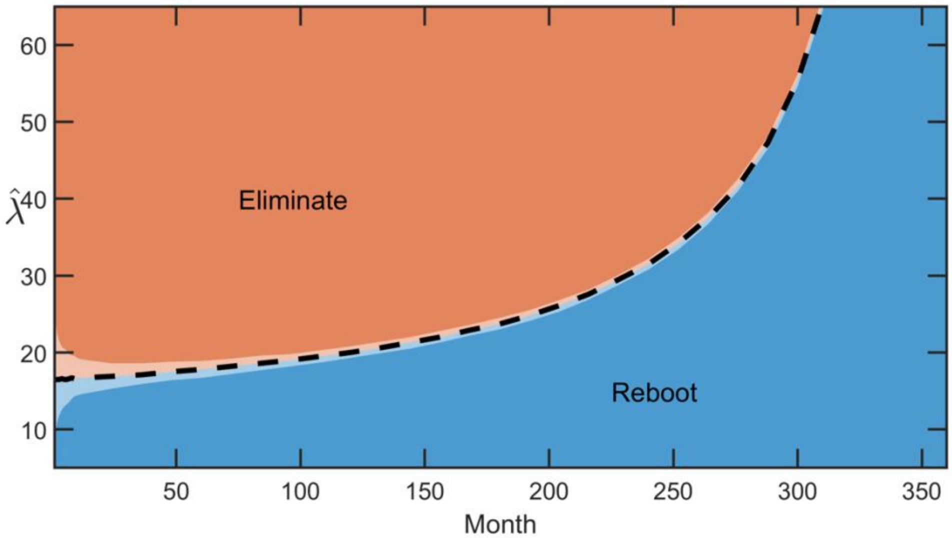

Figure 40: As we approach end of life, rebooting the system for the remainder of the life provides more utility than incurring high Eliminate costs. Only if the failure rate is very high, eliminating is worthwhile. $\hat{\lambda}_{threshold}$ *increases exponentially as we approach end of life*

The $\hat{\lambda}_{threshold}$ values shown above are for the case where Cost to Eliminate and Cost to Reboot are known for certain and the Monte Carlo simulation accounts only for the uncertainty in estimated failure rate. If we account for the uncertainty in costs as well, then the $\hat{\lambda}_{threshold}$ is higher (Figure 41). Therefore, it is worthwhile to Eliminate only at higher $\hat{\lambda}$ because there is added uncertainty. The cost uncertainties simulated here are:

$$\$800 < c_R < \$3{,}200$$

$$\$20{,}000 < c_E < \$228{,}000$$

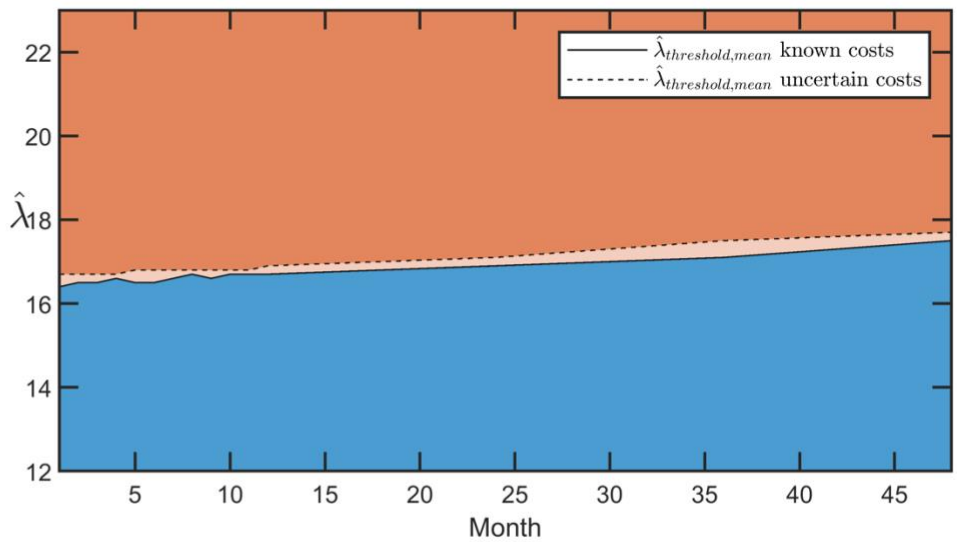

Figure 41: $\hat{\lambda}_{threshold}$ is higher when we account for uncertainty in costs

3.10 How does the decision framework help?

As laid out in the literature review, researchers say that organizations may not have an incentive to expend resources to find the root cause of an NFF. NFFs may not even be included in the MRO's failure statistics. The attitude towards NFFs is slowly changing and our research is to motivate the MRO organizations to take a closer look at the NFF problem and incentivize the need for eliminating or rebooting from a monetary/business perspective in terms of profit to the organization.

In Figure 42, we show a scenario where an existing MRO has been collecting failure data of an NFF for a long period of time. Since the number of failures vary every month, they calculate a running average of the failure rate and wait for the value to converge to start analyzing what action to take and whether it is worth investigating the root cause.

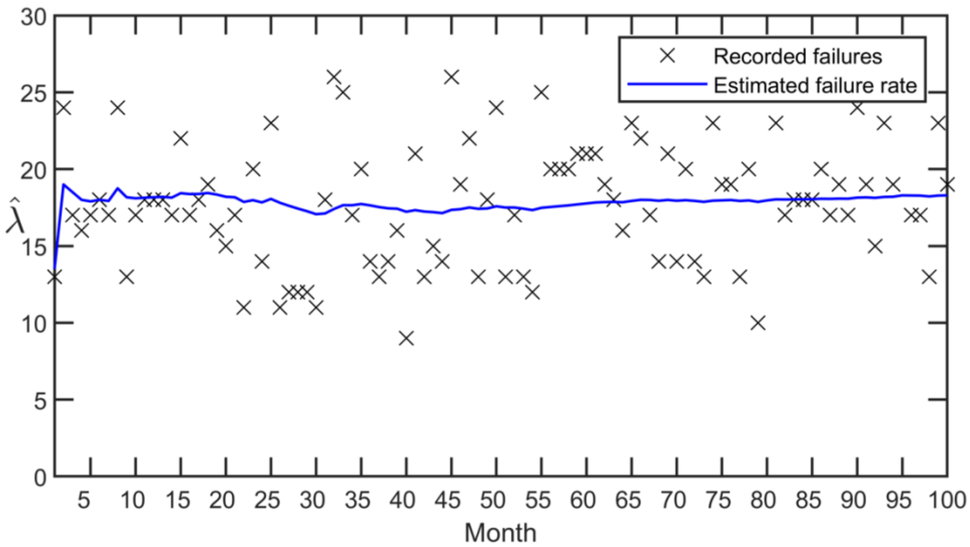

Figure 42: The MRO collects failure data of an NFF for a long time and waits for the estimate of the failure rate to converge to take any action

With our framework, we understand that waiting too long can lead to losing the opportunity to make the optimal decision. In Figure 43, we compare the estimated failure rate from the data collected by the MRO with the $\hat{\lambda}_{threshold}$. We see that, if you wait till month 100, then the $\hat{\lambda} < \hat{\lambda}_{threshold}$, which means Reboot is now more worthwhile.

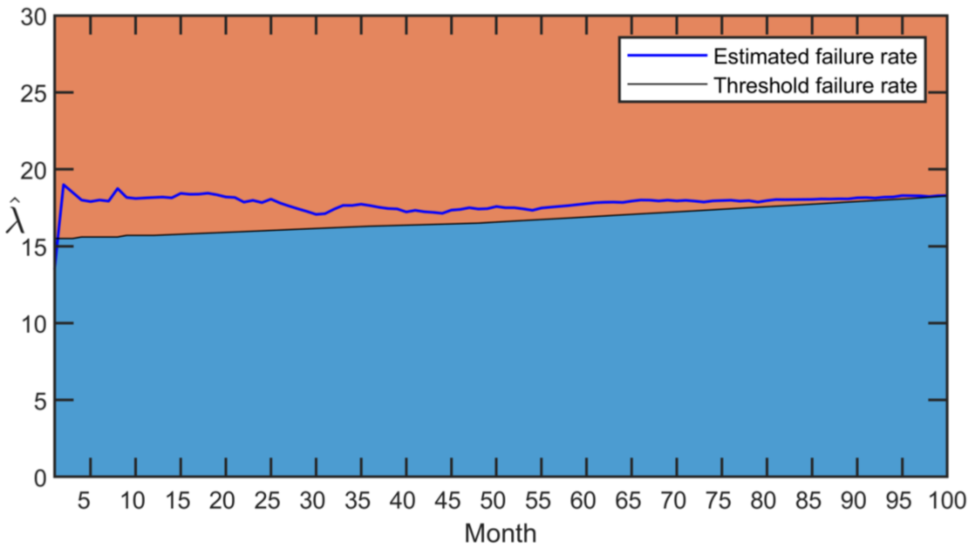

Figure 43:With the decision framework, we can compare the estimated failure rate with $\hat{\lambda}_{threshold}$ to find optimal decision. Here, if the DM waits for too long, they can lose the opportunity to Eliminate and make a profit

The converse is true as well. Making a hasty decision could also be undesirable due to uncertainty. The decision framework accounts for this uncertainty in information of cost and failure rates. In Figure 44, the uncertainty bounds for $\hat{\lambda}_{threshold}$ are also plotted. If the DM decides to eliminate early in the lifetime, without accounting for uncertainty, then there is a possibility of a loss. From the data in Figure 44, at month 20, with data available from the previous months, Eliminate is the optimal decision.

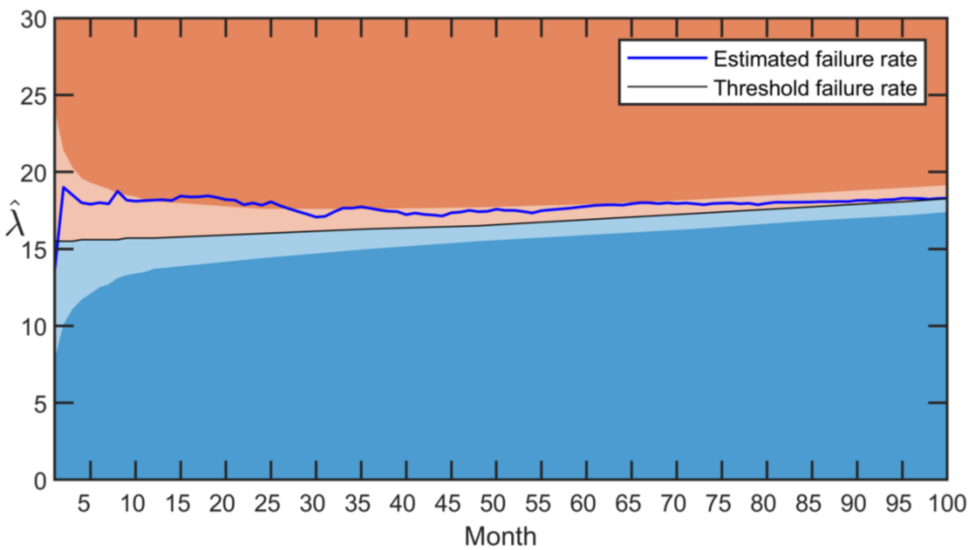

Figure 44: The decision framework accounts for uncertainty and therefore, at each decision point, the DM can know when it is better to wait for more information. In the case shown here, at month 5, the data is still uncertain and Reboot is optimal. Whereas at month 20, given the available data, Eliminate is the better decision

CONCLUSIONS

NFFs are a costly problem, both tangibly and intangibly. Further complicating this problem is the uncertainty in the occurrence rates of NFFs and the cost to recover from them. This work presents a Net Present Value and Expected Utility based analysis to aid in optimal decision making when faced with No Fault Found while accounting for these uncertainties.

The results presented are for the case of airlines that experience NFFs recurring due to a single failure mode in their fleet of a single aircraft type. Whether it is better to reboot or eliminate the NFF has a dynamism based on when the choice is made and how much information is available. There is value in waiting for more information and reducing uncertainty because the estimated failure rate might be far from the true underlying failure rate which would lead to a sub-optimal decision. Waiting for too long can also lead to losing the opportunity to make the optimal choice since it becomes less worthwhile to spend the large amount of eliminating as we approach the end of system life. In the initial phase of the system life, it is more worthwhile to eliminate the failure mode rather than continuing to reboot and incurring costs. As time progresses, we have less time to recover the cost of elimination and therefore it becomes more worthwhile to reboot.

We use a discounted cash flow method of NPV to quantify value provided by an aircraft fleet when experiencing recurring NFF events. The two decisions considered – Eliminate and Reboot – will give two NPVs. We model the uncertainty of failure rate (NFF occurrence rate) and uncertainty in the NFF costs which give a distribution of profit made when choosing to Eliminate. We employ Expected Utility Theory to model the risk attitude of the DM towards uncertain profits. We get Expected Utility of the Eliminate decision at each decision point which helps the DM choose between the two options.

We explore the change in behavior of NPV and Expected Utility with change in underlying failure rate, discount rate, risk attitude of the DM, Cost to Reboot and Cost to Eliminate. We find that NPV increases if discount rate is lower since value in the future is more worth in the present. We find that Expected Utility:

- Is more if underlying failure rate is higher since Eliminate decision leads to higher profit.

- Is more if DM is less risk averse and a loss is expected. Conversely, Expected Utility is lesser if DM is less risk averse and a Profit is expected.
- Increases with increasing Cost to Reboot and increases with decreasing Cost to Eliminate, since both these scenarios lead to more profits when deciding to Eliminate.

We introduce the concept of threshold failure rate estimate to provide the DM with a single parameter metric (given other inputs) to decide between Eliminate and Reboot. The DM can decide to make use of the lower bound, mean or upper bound of Expected Utility to derive the threshold failure rate. At the current decision point, if estimated failure rate for the future is greater than the threshold failure rate, then Eliminate is the optimal choice. If the estimated failure rate is less than the threshold value, then it is optimal to continue to Reboot.

We find that despite the uncertainties in the costs and NFF occurrence rates, a Utility maximizing choice between the two options considered – Reboot and Eliminate – is possible given the information available. It is important to consider the uncertainties to avoid hasty decisions but waiting too long for more certainty in data could lead to losing the opportunity to make the optimal choice.

In future work, some of these assumptions can be relaxed to expand the framework's fidelity and range of application, for example:

- Eliminating does not reduce the possibility of failure to zero. There is still some possibility of an NFF occurring in the component at a lower rate even after attempting to Eliminate. This scenario could be applicable in cases where, for example, the original NFF was found to be occurring because of human factors and addressing the root causes was to implement additional training for personnel. Here the training might be effective but may not eliminate the possibility of the fault occurring fully.
- Eliminate does not need to occur within the time between two decision points and can be uncertain (one month in our analysis).
- We can vary the underlying failure rate to model the component's deterioration over time.
- The framework can include the possibility of prior data being available before start of analysis or starting to collect data after some years of operation of the aircraft or fleet.

- The framework can include the possibility of multiple NFF processes acting at the same time.
- Different stakeholders can be added to the cost model (for example, External MRO, OEMs).
- The uniform distribution of NFF costs can be changed to a more realistic cost uncertainty distribution.

Obtaining actual data from MRO organizations would help validate the decision framework built in this research. Using this data, we can compare values of NPV calculated with the numbers estimated made by airlines MRO.

APPENDIX A. CONVERGENCE OF MONTE CARLO SIMULATIONS

In this work, we do 10,000 Monte Carlo runs to find Profit from deciding to Eliminate. In the Figure A.1, we show the convergence of the Monte Carlo runs using a running average of NPV_R calculated using baseline inputs (Table 1) when only the uncertainty in failure events is considered. The failures in the future months are Poisson distributed over the mean of estimated failure rate in month 1. The relative error, ϵ, as well as the average NPV_R, seem to have converged with 10,000 runs.

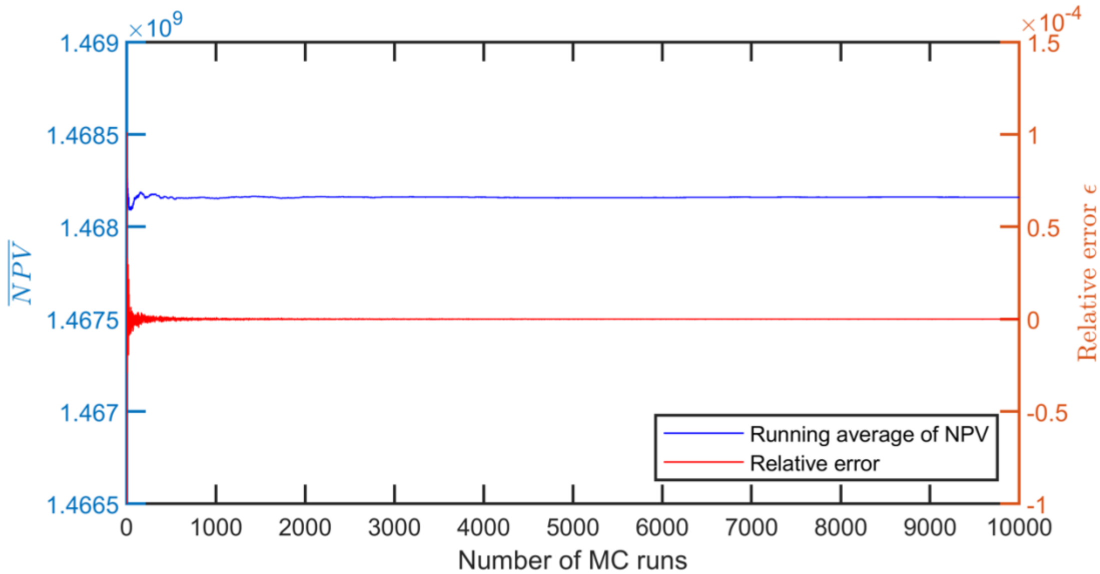

Figure A.1: NPV_R converges within 10,000 Monte Carlo runs. Here only the failure rate uncertainty is considered for the Monte Carlo simulations.

In Figure A.2, we show the convergence of the Monte Carlo runs using a running average of NPV_R calculated using baseline inputs (Table 1) when three uncertainties are considered, i.e., uncertainty in Reboot cost, Eliminate cost and number of failure events in the future months.

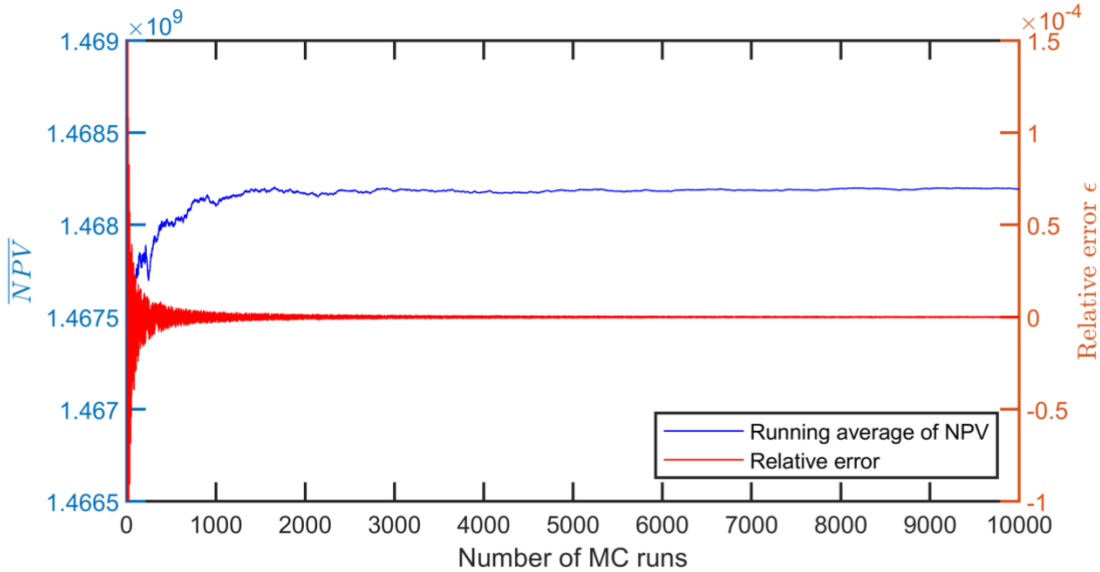

Figure A.2: NPV_R converges within 10,000 Monte Carlo runs. Here three random variables are used for the Monte Carlo simulations, i.e., the uncertainty in failure rate, Reboot and Eliminate costs.

REFERENCES

Air Transport Association of America [ATA]. (2002). ATA MSG-3 Operator/Manufacturer Scheduled Maintenance Development (Revision 2002.1). Air Transport Association of America, Inc.

ARINC 672. (2008). Guideline for the Reduction of No Fault Found.

Belobaba, P., Odoni, A. R., & Barnhart, C. (2009). The global airline industry. Wiley. https://onlinelibrary.wiley.com/book/10.1002/9780470744734

Beniaminy, I., & Joseph, D. (2002). Reducing the "No Fault Found" problem: Contributions from expert-system methods. Proceedings, IEEE Aerospace Conference, 6. https://doi.org/10.1109/AERO.2002.1036138

BTS. (2018). Average Age of Aircraft. https://www.bts.gov/content/average-age-aircraft

BTS. (2019). Airline Financial Data. First Quarter 2019 U.S. Airline Financial Data | Bureau of Transportation Statistics. https://www.bts.gov/newsroom/first-quarter-2019-us-airline-financial-data

Buede D. M., & Miller W. D. (2016). Decision Analysis for Design Trades. In The Engineering Design of Systems: Models and Methods (pp. 398–444). John Wiley & Sons, Incorporated.

Erkoyuncu, J. A., Khan, S., Hussain, S. M. F., & Roy, R. (2016). A framework to estimate the cost of No-Fault Found events. International Journal of Production Economics, 173, 207–222. https://doi.org/10.1016/j.ijpe.2015.12.013

Hölzel, N. B., Schilling, T., Neuheuser, T., Gollnick, V., & Lufthansa Technik, A. G. (2012). System analysis of prognostics and health management systems for future transport aircraft. 28th International Congress of the Aeronautical Sciences (ICAS), Brisbane, Australia.

James, I., Lumbard, D., Willis, I., & Goble, J. (2003). Investigating no fault found in the aerospace industry. Annual Reliability and Maintainability Symposium, 2003., 441–446. https://doi.org/10.1109/RAMS.2003.1182029

Kenley, C. R., & Armstead, D. C. (2004). Discounting models for long-term decision making. Systems Engineering, 7(1), 13–24.

Khan, S., Farnsworth, M., & Erkoyuncu, J. (2017). A novel approach for No Fault Found decision-making. CIRP Journal of Manufacturing Science and Technology, 17, 18–31. https://doi.org/10.1016/j.cirpj.2016.05.011

Khan, S., Phillips, P., Jennions, I., & Hockley, C. (2014). No Fault Found events in maintenance engineering Part 1: Current trends, implications and organizational practices. Reliability Engineering & System Safety, 123, 183–195. https://doi.org/10.1016/j.ress.2013.11.003

Kinnison, H. A., & Siddiqui, T. (2013). Aviation maintenance management (2nd ed). McGraw-Hill.

Kumar, U. D., Crocker, J., Chitra, T., & Saranga, H. (2006). Reliability and Six Sigma.

Marais, K. B. (2013). Value maximizing maintenance policies under general repair. Reliability Engineering & System Safety, 119, 76–87. https://doi.org/10.1016/j.ress.2013.05.015

Marais, K. B., & Saleh, J. H. (2009). Beyond its cost, the value of maintenance: An analytical framework for capturing its net present value. Reliability Engineering & System Safety, 94(2), 644–657. https://doi.org/10.1016/j.ress.2008.07.004

MIT. (2018). Airlines Data Project: United Airlines. http://web.mit.edu/airlinedata/www/United.html

Montgomery, D. C. (2005). Introduction to statistical quality control (5th ed). John Wiley.

Qi, H., Ganesan, S., & Pecht, M. (2008). No-fault-found and intermittent failures in electronic products. Microelectronics Reliability, 48(5), 663–674. https://doi.org/10.1016/j.microrel.2008.02.003

Raza, A., & Ulanskyi, V. (2016). Assessing the impact of intermittent failures on the cost of digital avionics' maintenance. 2016 IEEE Aerospace Conference, 1–16. https://doi.org/10.1109/AERO.2016.7500563

Sahai, H., & Khurshid, A. (1993). Confidence Intervals for the Mean of a Poisson Distribution: A Review. Biometrical Journal, 35(7), 857–867. https://doi.org/10.1002/bimj.4710350716

Sahay, A. (2012). Leveraging information technology for optimal aircraft maintenance, repair and overhaul (MRO). WP, Woodhead Publ.

Saleh, J. H., & Marais, K. (2006). Reliability: How much is it worth? Beyond its estimation or prediction, the (net) present value of reliability. Reliability Engineering & System Safety, 91(6), 665–673. https://doi.org/10.1016/j.ress.2005.05.007

Söderholm, P. (2007). A system view of the No Fault Found (NFF) phenomenon. Reliability Engineering & System Safety, 92(1), 1–14. https://doi.org/10.1016/j.ress.2005.11.004

Sudolsky, M. D. (1998). The fault recording and reporting method. 1998 IEEE AUTOTESTCON Proceedings. IEEE Systems Readiness Technology Conference. Test Technology for the 21st Century (Cat. No.98CH36179), 429–437. https://doi.org/10.1109/AUTEST.1998.713480

Swanson, L. (2001). Linking maintenance strategies to performance. International Journal of Production Economics, 70(3), 237–244. https://doi.org/10.1016/S0925-5273(00)00067-0

Thomas, S., B, H. N., & Stephan, L. (2012). Cost-Benefit Evaluation of Aircraft Maintenance Base-Check Downtime and Cost. Stand Alone, 371–378. https://doi.org/10.3233/978-1-61499-119-9-371

Werner, D. (2015). A maddening costly problem.(no fault found)(Cover story). Aerospace America, 53(2), 28–33.