Application of Multiple Information Sources to Prediction of Engine Time On-wing

by

Daniel Richard Roberson

B.S. Physics, University of Portland, 2006

Submitted to the MIT Sloan School of Management and the Department of Aeronautics & Astronautics in Partial Fulfillment of the Requirements for the Degrees of

Master of Business Administration
and
Master of Science in Aeronautics & Astronautics

In conjunction with the Leaders for Global Operations Program at the
Massachusetts Institute of Technology
June 2015

Signature of Author__________Signature redacted_
MIT Sloan School of Management
Department of Aeronautics & Astronautics
May 8, 2015

Certified by________Signature redacted________
Dr. Youssef Marzouk,
Thesis Supervisor, Associate Professor of Aeronautics and Astronautics
Director of Aerospace Computational Design Laboratory

Certified by________Signature redacted________
Dr. Karen Willcox,
Thesis Supervisor, Professor of Aeronautics and Astronautics
Co-Director of Center for Computational Engineering

Certified by__________Signature redacted_
Dr. Roy E. Welsch,
Thesis Supervisor, Professor of Statistics and Management Science

Accepted by________Signature redacted________
Dr. Paulo Lozano, Chair of the Graduate Program Committee
Associate Professor of Aeronautics and Astronautics

Accepted by __________Signature redacted__
Maura Herson,
Director of MBA Program MIT Sloan School of Management

This page has been intentionally left blank.

Application of Multiple Information Sources to Prediction of Engine Time On-wing
by
Daniel Richard Roberson

Submitted to the MIT Sloan School of Management and the Department of Aeronautics & Astronautics on May 8, 2015 in Partial Fulfillment of the Requirements for the Degrees of Master of Business Administration and Master of Science in Aeronautics & Astronautics.

Abstract

The maintenance and operation of commercial turbofan engines relies upon an understanding of the factors which contribute to engine degradation from the operational mission, environment and maintenance procedures. A multiple information source system is developed using the Pratt & Whitney engine to combine predictive engineering simulations with socio-technical effects and environmental factors for an improved predictive system for engine time on-wing. The system establishes an airport severity factor for all operating airports based upon mission parameters and environmental parameters. The final system involves three hierarchical layers: a 1-D engineering simulation; a parametric survival study; and a logistic regression study. Each of these layers is combined so that the output of the prior becomes the input of the next model. The combined system demonstrates an improvement in current practices at a fleet level from an R^2 of 0.526 to 0.7966 and provides an indication of the relationship suspended particulate matter and engine degradation. The potential effects on the airline industry from city based severity in maintenance contracts are explored. Application of multiple information sources requires both knowledge of the system, and access to the data. The organizational structure of a data analytics organization is described; an architecture for integration of this team within an existing corporate environment is proposed.

Thesis Supervisor: Dr. Youssef Marzouk
Associate Professor of Aeronautics and Astronautics
Director of Aerospace Computational Design Laboratory

Thesis Supervisor: Dr. Karen Willcox
Professor of Aeronautics and Astronautics
Co-Director of Center for Computational Engineering

Thesis Supervisor: Dr. Roy E. Welsch
Professor of Statistics and Management Science

This page has been intentionally left blank.

Table of Contents

List of Figures

Glossary of Terms

A1, A2, A3, A6, A8, A14, A19, A21 – Aerosol volume column loading parameters defined in Table 2
ACARS - Aircraft Communication Addressing and Reporting System
AIC - Akaike information criterion
ASF – Airport Severity Factor
BIC - Bayesian information criterion
Business unit – team within a company division responsible for specific goals
CALF – Configuration adjusted life factor
Derate – The setting below maximum available thrust, measured as a percentage where 0 equals max thrust and 100 equals zero thrust
DTamb – Difference between the ambient temperature and the elevation adjusted temperature at International Standard Atmosphere
EGT – Exhaust Gas Temperature
EPR – Engine Pressure Ratio
ES – Engine Services, a business unit within P&W aftermarket division
FHA – Flight Hour Agreement
GE – General Electric Corporation
GW – Gross Weight at takeoff
HPT – High Pressure Turbine

HPT Grids – Data set provided by Hot Section Engineering (HSE) that describes the relationship between expected Time on Wing (TOW) of the High Pressure Turbine (HPT) and operational parameters of the mission point
HSE – Hot Section Engineering, a business unit within the P&W Engineering division
IAT – Interval Analysis Tool
IAE – International Aero Engines
Interval – Time between hot section refurbishments (HSR) measured in either flight hours or flight cycles
ISA - International Standard Atmosphere
IT – Information Technology, in reference to an organization within a company
JMP – A statistical software from SAS Corporation
LF – Life Factor
LLP – Life limited part, parts that require replacement after fixed cycles or hours regardless of condition
MISR - Multi-angle Imaging SpectroRadiometer
Mission point – a characteristic mission of an engine described by all pertinent characteristics of a flight route between two typical airports
MODIS - Moderate-Resolution Imaging Spectroradiometer
MRO – Maintenance and Repair Organization
MSE – Mean square error
OAG – Official Airline Guide
OEM – Original Equipment Manufacturer
P&L – A segment of a company measured against for profit and loss metrics for all activities
PM – Particle Mass, measured in column loading at or above a specified particle size
P&W - Pratt & Whitney, a UTC company
RH – Relative humidity
Severity – A factor describing the difference between expected life and actual life as a function of unknown parameters. Larger numbers correspond to longer life.
SPSS – Statistical Package for the Social Science, a software system from IBM Corporation
SQL – Structured query language
TOW – Time on wing, typically measured in flight hours
TR – Thrust Reverser
UTC – United Technologies Corporation

1.0 Introduction

As the growth of data systems has made Big Data a buzz word in industry, organizations struggle to develop methods capable of managing large data systems and harvesting financial value from them. Two problems are specifically addressed in this work: the problem of properly maintaining and propagating big data through a company, and the problem of developing successful models for multi-factored socio-technical scenarios. The description of state of the art model applications within the aerospace and insurance industry provides context for the current work. The specific area of research is the advancement of engine time on wing (TOW) prediction, subject to variability from basic engine physics, maintenance and inspection policy, and airport specific environmental effects. This chapter introduces the current state and objectives of engine TOW prediction at Pratt & Whitney, and outlines the methods by which a new model is proposed.

1.1 Background

1.1.1 Applications of Models Across Industries

Analytical models throughout multiple industries demonstrate similarities in the way they merge diverse data sources to generate complex structure simulations. Both the aerospace maintenance and health insurance industries insure the risk of expensive rare events and take advantage of similar cost modeling methods. Health insurance decision models have been developed using Monte Carlo simulation[1], Markov Chain analysis[2], data clustering[3], spatial geographic[4], mixed financial steady state models[i] and neural networks[5]. The varied application of data modeling tools indicates that these industries hold sufficient data to answer questions in several ways.

Often the model selection and the type of question being asked of the data are decided together. Many of these models in healthcare take advantage of mixed data sources and mixed model methods when questions require it. For example, a first order model found in health care is the Proportional Disease Magnitude (PDM), which is an intermediate factor that represents the financial cost of a disease in the place of actual costs. This PDM is first developed for an area of interest, and then used as a factor in simulations[6],[7]. The abstraction of this information provides greater information for portfolio level risk analysis than the unadjusted cost data would have done since the underlying factors of cost can be separated from the noise incurred by billing errors.

1.1.2 Engine Support Policies and Predictive Models

Airplane engine maintenance represents 41% of the cost of ongoing maintenance for worldwide passenger airline operators [8]. For both airlines and engine manufacturers, the decisions regarding maintenance costs management are a critical aspect of the original purchase decision. Engine maintenance contracts are awarded competitively by the airplane owner to one of three options: Independent Maintenance Repair & Overhaul (MRO) companies, Airline Operated MRO's, and the engine's primary original equipment manufacturer (OEM). From the airline operators' perspective, these contracts serve to mitigate risk by either setting repair rates or directly covering certain repairs. Since 1962, with Rolls Royce's *Power By-the-Hour*™ plans, many operators pay fixed amounts per flight hour to the OEM for the privilege of zero cost, or decreased cost shop visits[9]. Referred to as Flight Hour Agreements (FHA), these fixed rate contracts are billed either monthly or at shop visit. The coverage of such plans is typically limited to a fleet cumulative term, per engine term, or a fixed number of shop visits. Implications of each of these combinations of billing and coverage methods complicate the pricing decisions for both the operator and the OEM. The actual cost of operating the fleet is subject to a high degree of uncertainty with high costs incurred infrequently. The worldwide average maintenance expense in 2011 was just over $200 per flight hour, and engines perform on the order of 10,000 flight hours across a time horizon of several years before overhaul costs are first incurred[8]. The industry standard method of modeling this time on-wing risk is to use Weibull distribution analysis based upon existing data, or engineering prediction[10]. Simulation and risk based customer pricing are developed by considering a number of factors including operating thrust, environmental conditions, flight length, and engine configuration[11]. Accurate understanding of the expected time on-wing (TOW) under a set of operating conditions is central to developing a cost structure for FHAs. Understanding the expected distribution profile for TOW enables accurate simulation of the differences between term limited plans and shop visit based plans and optimization of fleet management strategies.

The day to day operating conditions of an engine are collectively referred to as the mission. Mission parameters include such values as the thrust, hours per cycle, and local operating conditions. In addition to mission parameters, other engine lifelong or fleet level factors affect TOW including maintenance policy, thrust reverser usage, fuel quality, and others. While FHA pricing takes into consideration all of these factors, it is difficult to accurately predict the expected life of a new engine program across all possible missions. First principles simulation of the engine gas path is performed by the engineering team for a variety of inputs and design modifications are tested against prior versions to mitigate against key failure modes identified in TOW predictions. Taking into consideration all possible combinations of local operating conditions increases simulation cost and any complete model of all interactions would require re-certification on each engine program. As a result, the industry commonly takes advantage of TOW adjustment tables that are treated as independent or sometimes covariate effects. To generate a predicted TOW, the baseline prediction from engineering design is obtained and multiplied by a set of severity factors from these tables. A severity factor of 1.0 is used as the base point for calibration of the first principles models and the real world experiences on prior engines. By convention high severity indicates proportionally poor performance. For example, thrust derate is measured as the percent below maximum rated thrust. It is selected by the pilot at takeoff. Raising thrust from a derate of 10% to a derate of 0% may change the severity factor from 1.0 to 1.1 yielding a 10% shorter expected TOW[11]. These tables are developed by the manufacturer and provided to prospective airlines when planning their mission profiles and maintenance costs. Many mission parameters are directly related to one another as will be described. In these cases, a table or set of tables showing the interacting terms is constructed, while for independent effects a collection of tables can be applied sequentially. The use of severity factors as a process for predicting engine TOW enables the manufacturer to solve some problems independently and enables the transfer of information forward to future models.

1.1.3 The Objective - City Level Severity

Geographic dispersion of the airline industry is itself a recent trend. In 1995 the regions of North America, Europe, Pacific Islands or Seaboard garnered 84% of all passenger traffic. Dusty regions, such as North Africa, the Middle East, South Central Asia and inland China did not reach 10% of global traffic until the year 2001. Between 2005 and 2010 the air demand in these latter regions doubled to a half billion passengers per year [12]. Several years later, the first engines began returning for repairs in high enough numbers for statistical analysis. The impacts of dusty environments on fleet level TOW had therefore not been widely studied until 2010. Recent work by multiple manufacturers focuses on the development of new severity factors to account for regional effects. The relationships between particle type and size and the associated engine degradation modes are being explored. As a result of these studies, one manufacturer has stated that the "TOW ... is a function of percent of routes in harsh environments"[13].

The severity factor of an airport environment could be estimated by first principles simulation through a full characterization of the types of particles in the operating space. This analysis could culminate in proper understanding of the material science and generate useful results for the development of new engines and simulation of existing fleet performance. The results would not be readily absorbed by the industry or fleet operator and would be expensive to generate and to run. By contrast, the development of a severity factor that accounts for the operational environment by regression to globally available values could be easily adopted by the manufacturer and the operator for fleet planning.

It is the objective of this study to generate an airport severity factor (ASF) for use in calculating a fleet level severity factor. By combining ASF with other effects of the mission including flight length, it will then be possible to produce city pair level estimates. City pair estimates are the target of the industry according to Jim Pennito, "We're trying to understand whether a certain city pair is more abusive to the engines than another city pair[14]." To be applicable in sequence with the existing set of severity factors, the properties of the ASF must exhibit substantial independence from existing severity factors and engine lifetime predictors. The correlation between the driving factors of the ASF and existing severity factor variables will be explored.

1.2 Methods

The application of multiple information sources method is capable of providing near term benefit that cannot be achieved by complete single method models in the context of an actively used organizational process. Multiple information source methods take advantage of a variety of analytical modeling and decision making processes in a hierarchical order[15],[16],[17]. The objective of these methods is to abstract key information about the behavior of a complex system which will tend not to behave in closed form manner. Multiple information source systems mimic the way people and organizations make decisions. Certain stages of the decision process involve parameters that are excluded from other stages. Such decision making systems can be applied to any existing model structure to improve information captured by the organization in an incremental way by building upon current processes and data.

The development of a multiple information source system requires a full understanding of the data pertinent to the decision or set of decisions which is to be made. The context of the decision and maintenance of the system are critical aspects of organizational intelligence that contribute to the value, flexibility, and relatability of the final product. The target variable of interest to this study is engine time on-wing (TOW). The decision being made is a service contract pricing decision. This study starts with a system dynamics approach to capture the behavior of both the pricing decision and the engine TOW. Physics effects that require special treatment are identified. These effects inform the selection of cross terms applied at individual layers of the model and the development of derived variables for model augmentation. Collinear terms are considered prior to model development to prevent improper training of the model. The data sources are identified using the system dynamics model along with the level of data validation and re-formatting. Reformatting improves the value for particular applications and may lead to information loss in others. Finally, data aggregation is performed to improve response specificity and correct for sampling biases incurred by missing data. The treatment of the data depends upon the question being asked. Figure 1 shows the proposed system of a three layer model which isolates predictable physics effects from engine removal conditions and airport severity.

Figure 1: Proposed Multiple Information Source Model for TOW prediction

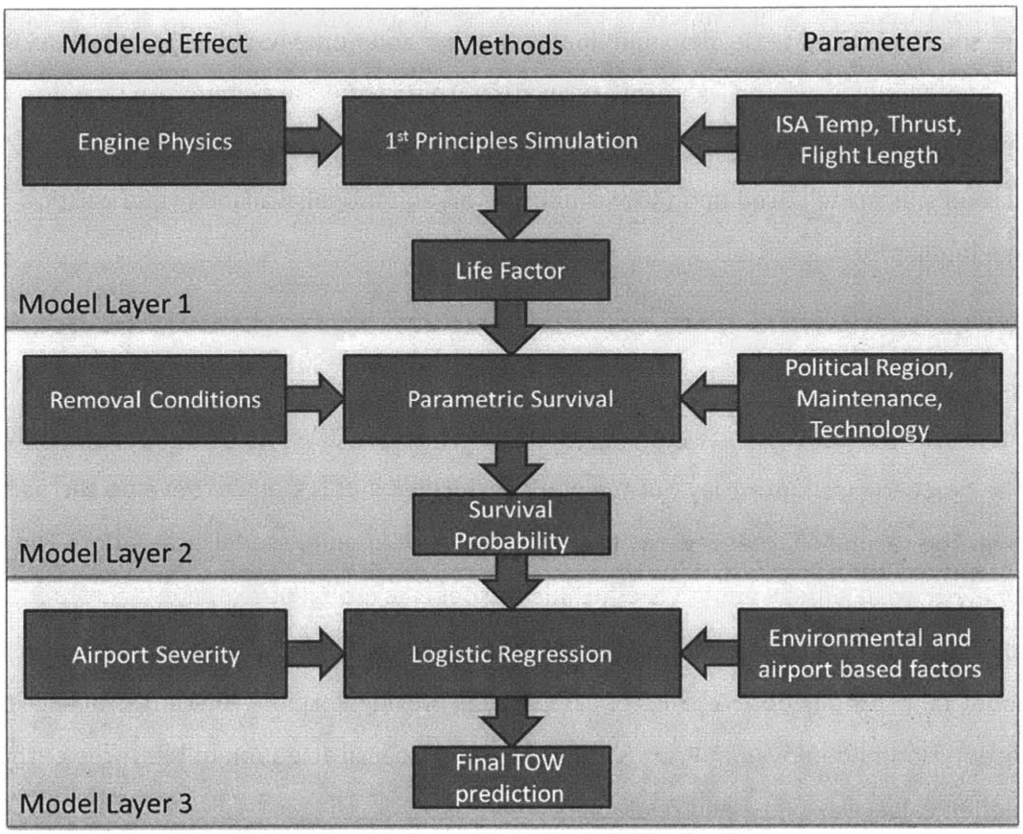

This study applies system dynamics modeling as a first step toward defining interaction terms for use. This system dynamics model also informs the process in model development of obtaining data in cases where not all factors can be obtained easily. Next the known interactions of operating parameters within the engine are obtained. Physics is used in three ways to analyze the problem structure before performing data gathering and effect screening: 1) units and distribution effects, 2) isolation of factors, 3) data augmentation by new variables.

The two primary data sources of the model are Aircraft Communication Addressing and Reporting System (ACARS) and NASA MISR. A data rationalization policy is established for each variable that converts the variable into proper unit basis as determined by the physics analysis, and null records are either filled with reasonable estimates or left blank.

Aggregation of data prior to or concurrent with modeling is performed at several layers. Aggregation of rows in datasets improves the model specificity and can be used to improve signal-to-noise, both of which improve model training accuracy. In each aggregation layer two important factors are considered: Sampling effects of aggregation basis; Information loss on non-linear effects.

Four software suites are used to develop and deploy the final solution: SAS JMP 11.0, Matlab 2014a, Teradata 14; SPSS Modeler 16; and Microsoft Excel 2010. These software systems are used for data pre-processing, enterprise data integration, model development and model execution respectively. The software decision was made based upon tool capability, organizational availability and existing organizational familiarity.

1.3 Summary of Chapters

The objective of this study is to develop and deploy an improved method of handling airport level differences in engine performance in such a way that the effect on TOW can be more accurately predicted. The proceeding chapters lay out the work performed in this study to develop the particular solution above. The second chapter reviews the state of the art in multi-model integration. The literature review gives attention to both the technical aspects of model development across the variety of fields involved in this particular problem and the management aspects of data maintenance – a key aspect of model value proposition. A series of case studies provided explore the causes of success or failure for recent models deployed at Pratt & Whitney giving special attention to the three criteria for a successful model: value, flexibility, and relatability.

Chapter three details the work performed in establishing the entire solution. A study of the system dynamics involved informs the selection of initial factors for the decision structure. Next we calculate the physics interactions of higher order effects predicted by the system dynamics model. An explanation of the proficiency and predictive nature of a variety of modeling methods validates the selection of the models used in the final system. Finally the results of each layer of the final system are displayed and the sensitivity of these models to underlying data assumptions and methods is examined.

Chapter four performs validation of the final airport severity factors against internal benchmarks to Pratt & Whitney through comparison with existing data models for regional severity and by examining the sensitivity of the factors to variance in source parameters. The organization supporting and implementing this system requires analysis to ensure proper deployment is sustainable and that value, flexibility, and relatability are communicated to the company.

Chapter five reviews the impact that the airport severity factors will have upon the existing organization, both internally to Pratt & Whitney and to its customers. The relatability of the ASFs makes possible the expansion of them into future market contracts across the industry and the potential effects of this shift are discussed. Chapter five also covers the logical extensions of this work both in the near term range for other applications of multiple information source methods and for evaluation of new IT and data management policies.

2.0 State of the Art

Extensive literature on the topic of engine physics and large data systems provides a background for the current research. The combination of four areas of literature is integral to establishing a functional model that is sufficiently complex to handle the effects present in the lifetime of an engine. Not only does this system of models require an understanding of engine physics and data modeling approaches, but the organizational methods of a company require careful examination to ensure support exists for the data system. This chapter reviews the research in system dynamics, engine physics, multiple information source systems, and organizational structures for the support of such systems. Two case studies from Pratt & Whitney are discussed which highlight shortcomings of two data models and suggest how these models could have been improved by management changes during their development and deployment. Finally, the current state of TOW modeling is described within the context of the P&W data systems.

2.1 External Literature Review

2.1.1 System Dynamics

System dynamics in its core describes a method of capturing interactions present in a complex system with or without given mathematical relationships. System dynamic modeling in reliability engineering tends to be done in two high level categories: Markovian decomposition, and multi-parameter feedback analysis. Markovian decomposition of the system enables a variety of modeling methods for both error classification techniques [18], [19] and system level reliability [20], [21]. In each of these cited examples, the system dynamics diagram informs the hierarchy of the ensuring model by capturing the expected behavior and interactions prior to mathematical screening. Chung et al develop a comprehensive dynamics model of module-to-module interactions within turbofan gas paths which enables the construction of a complex system level model using simple linear interactions between each pair of nodes[22]. System dynamics is applied to analytical methods outside aerospace to effectively model socio-technical characteristics alongside empirical mathematical methods, most recently in the healthcare and insurance industry [23], [5]. One example of hybrid modeling including regulatory policy evaluation and stochastic mathematic modeling of flood risk developed for European drainage simulations relies upon system dynamics to organize influential factors [24].

2.1.2 Engine Physics

Estimation of internal engine temperatures and pressures is a central aspect of the design phase for both production and market placement. The basic methods and trade-offs are published by both manufacturers and independent parties [25], [26], [27], [28], [29]. While OEMs use proprietary modeling methods for design and evaluation, academic versions of models such as Hermes and Turbomatch developed at Cranfield University for airframe and engine performance modeling respectively provide sensitivity analysis of engine durability to standard factors. Modeling of the internal temperatures and pressures requires assumed values for external atmospheric values throughout the flight, and a full thermodynamic model of the engine in question. The majority of modern methods approved within industry for performing the gas path simulation involve recursive linear quadratic estimation known as a Kalman filter[30],[31], [32]. This method, based on an underlying Bayesian inference structure with Gaussian assumptions, is capable of developing a stable solution to internal temperatures and pressures given limited known constraints and physics with significant noise. Innovative methods have been presented to more closely simulate engine health performance using non-linear and non-Gaussian methods with Monte Carlo sampling methods to constrain solutions on an extended non-linear Kalman filter with non-Gaussian inputs[33]. This work has shown a high degree of sensitivity to methods of sampling. While these methods may eventually be adopted for future systems, the commonly deployed models in industry remain linear and Gaussian in nature.

Sensitivity analysis between these mission parameters and engine life expectancy is based upon simulation of low cycle metal fatigue, material oxidization, and high cycle creep through time at temperature estimations. Hanumanthan provides a thorough review of the methods for combining damage estimators to predict engine failure [34]. The results of these models enable declarable trade-offs specific to the engine model such as "A takeoff derate that averages 5% will add 400-500 [engine flight cycles] to on-wing life, a 10% derate will add 800-900, and a 15% derate will add about 1100[35]."

The basic physics of engine propulsion begins with the equation of net thrust:

Equation 1 $$F = \frac{W_a}{g}\left(V_j - V_a\right) + A_j(P_j - P_{am})$$

where W_a is the air mass flowrate, g acceleration of gravity, V_j the exhaust average velocity after mixing in the exhaust cone, V_a the incoming air velocity on the axis of motion, A_j the aera of the nozzle, P_j the static pressure at the discharge and P_{am} the static pressure of ambient air[28]. In the above formulation, momentum of spent fuel (W_f) is disregarded. From this formula and basic atmospheric relationships, the impact of temperature, pressure, and wind speed on the gross thrust are readily established. Turbines follow a Brayton cycle, shown in Figure 2, capturing energy on the expansion path to power the compressor, the aircraft electronics and the thrust.

Figure 2: Gas Turbine Work Cycle

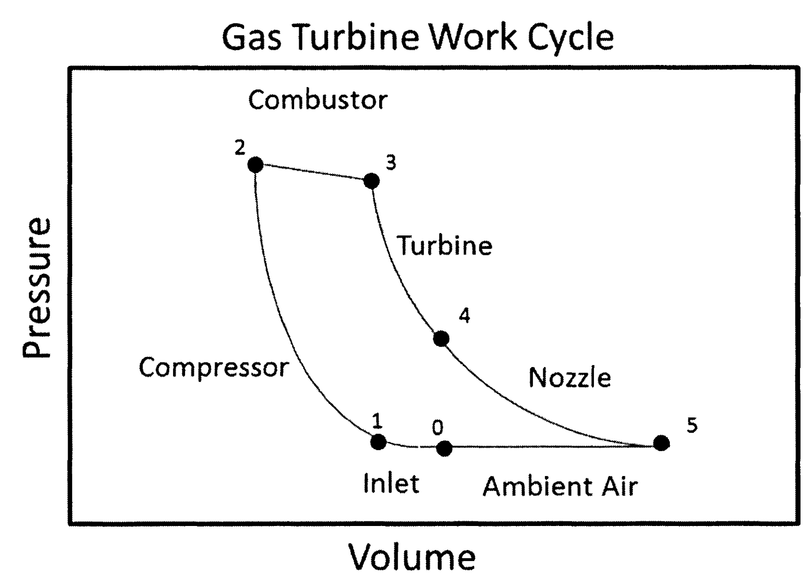

Since fixed nozzle designs are commonly used for high by-pass systems used on standard commercial aircraft applications, the expansion curve does not change to adjust for flight conditions. Until the introduction of the Pratt & Whitney Geared Turbo Fan, the relationship between compressor and expansion curves was fixed by the common spindle[36]. With fixed engine geometry, any change in thrust caused by atmospheric properties may only be countered by increasing combustor fuel flow and consequently internal gas path temperatures. With atmospheric values as constants, the engine thrust can be rewritten as $F = C_{atm} + A_j * EPR$ where C_{atm} collects terms fixed by geometry, airspeed and ambient pressure and EPR is the engine pressure ratio defined as the pressure at the nozzle divided by the pressure at the compressor: point 4 over point 2 in Figure 2. This formulation of Equation 1 informs a number of severity assessments in that the EPR, thrust and fuel flow are shown to be linearly related at any given operational point.

Atmospheric and mission parameters are highly covariate both through standard atmospheric equations and through operational constraints. For example, Mishra and Beura demonstrated that the relationship between runway length for maximum load aircraft and thrust requires linearly higher thrust for shorter runways and explored the effect of this and other factors on the engine life consumption[37].

2.1.3 Multiple Information Source Learning Methods

Multiple layer models have been developed to perform remaining life estimates on engines where single performance parameters such as Exhaust Gas Temperature may be used as suitable evaluations of health[18]. The method developed by Ramasso and Gouriveau first performs a prediction of the failure mode which will occur using an evolving neuro-fuzzy system, next the engine is classified into a time-risk based mode using a parametric Markovian classifier. The model performs automatic clustering enabling the second layer to build time series categories without prior knowledge except a basic belief assignment which establishes the underlying expected variance leading to a failure. While the later aspect of this approach is studied independently [19], the combined model improves both the accuracy and flexibility of the system.

The Climate Risk Insurance Model (CRIM) represents an advanced use of multiple information source modeling[4]. The CRIM could not be developed using a single equation or modeling method and requires the use of nested model layers. In each layer, the modeling method is selected based upon the information type. Climate Change scenarios provide inputs to a region specific flood probability in Netherlands dikes. These models are combined with land use scenarios based upon socio-economic development models to generate a damage function per geography, flood depth, and time. While economic development scenarios may be modeled with regression techniques, the risk models of the combined data are suited to Monte Carlo simulation. The model performs simulation on the combined cost model for 100 years of future climate and development inputs. The output probabilities are converted to a risk premium against a number of insurance policy structures to evaluate both the insurance pricing and insurance liabilities. Finally, these results generate price sensitivities for the Netherlands in regards to the development and maintenance of new dikes.

2.1.4 Organizational Structures for Data Integration

Data storage and curation decisions impact the owning company's ability to make productive use of the data for decades after the decisions are made. Decisions about data management fall into three main categories: Storage medium, Database architecture, Data curation[38]. Storage medium describes the physical storage location and includes decisions about physical security, backup reliability, geographic emergency preparedness, and power systems. Database architecture refers to the software deployed on the storage medium. These decisions involve trade-offs between storage volume, speed and accessibility. The database architecture decision cannot be made without an understanding of the nature of the data to be stored. After establishing these two underpinnings of infrastructure, data curation refers to all activities of data management from acquisition to delivery including verification of usability, applicability and advertising of the data[39].

While many models exist for data management, there is only one Internationally approved standard which provides a strong baseline for discussion: the Open Archival Information System (OAIS)[40] [41]. The OAIS, shown in Figure 3, describes data curation in six areas. Main data flows are connected by solid arrows with the data format of the transfer being called out in the nearby ovals. Functional oversight is connected by dashed lines. Descriptive Info refers to the metadata and schema definitions required to define the data contents and storage types. SIP, AIP and DIP refer to Submission, Archival, and Dissemination Information Packages respectively and have defined standards for metadata included in each package. Ingest describes those processes controlling the collection of new data into the data system. Following Ingest, data rationalization and storage to the servers is governed by the Archival Storage function. In parallel Data Management functions perform documentation and integration of the data to the overall system architecture. The final inflow function of Access works with consumers to identify the existence, description, location, and availability of data. Over the daily processing of data the Preservation Planning function maintains hardware and software general operations, working with consumers and producers to develop and maintain availability requirements and update plans. Finally, the Administrative function provides integration of all areas of the archive system and is responsible for those decisions affecting multiple functions.

Figure 3: OAIS Data Curation Model

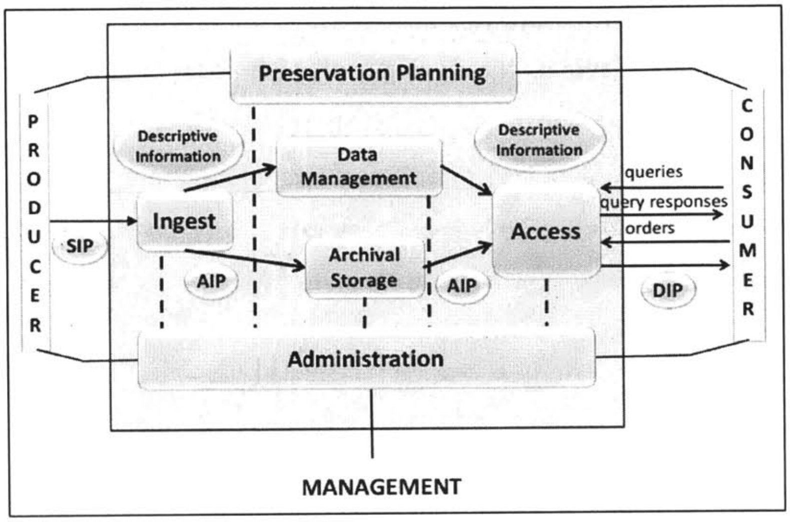

Data Integration within a corporate environment is challenged to provide both large data storage and simultaneously well exposed data. Challenges are posed by mixed data formats between both digital data and traditional document based analogue data such as receipts or maintenance notes. Activity based pricing has shown that Preservation Planning, Administration, and Archive Storage are the most expensive activities of a data system[42, p. -]. This cost centering generates risk of a myopic view of data curation. Studies performed on the World Data System proposed the value of a formally defined pre-ingest function and the development of mirrored processes for dealing with analogue and digital data in a holistic manner through the Data Management function[43]. A survey performed on 26 members of the WDS established a scoring method to determine the effectiveness of each function. This survey revealed that Ingest and Data Management are generally the strongest parts of a data system with results by group shown in Table 1. Access is the weakest on average with a number of very low scoring systems.

Table 1: Survey Scores for WDS functions, adapted from [44]

	Ingest	Archive	Data	Admin	Preserve	Access
Average Score	76%	66%	79%	65%	63%	54%
1 Sigma Range	53%-98%	49%-85%	61%-95%	48%-81%	40%-86%	27%-81%

2.2 Internal Case Studies

We propose that effective models in a business context require three principal components: value, flexibility, and relatability. Value refers collectively to the strength of the model, the validity of the business need it fills and the cost of running it. Flexibility encompasses the range of unpredictable inputs the model must allow for while assisting in the business process. Relatability requires that the customers and operators of the model have a proper understanding of how the system works, which enables them to make full use of the information provided. This section presents two unique cases observed in which a model or system of models fell short of attaining one of these three elements. The following cases do not represent failed projects. They provided value and contributed to company success. Rather they represent projects which fell short of their potential. Each case study concludes with recommendations for ongoing improvement.

2.2.1 Event Prediction Models

2.2.1.1 Objective Description

Teams of engineers review performance data of engines daily through a proprietary on-wing and off-wing data management system, monitoring fleets for any indication of reliability concerns and responding quickly to operational events. Engine deterioration drives scheduling of inspections and engine overhauls with the majority of engines never having any events in their on-wing lifetimes. In this setting, the company pursued the development of a predictive model that could identify even minor events before they occur.

2.2.1.2 Model Development

The development team reviewed the current process for event detection and developed a list of potentially interesting sources of data. Over several weeks a variety of data was acquired from multiple information sources and rationalized together. It was observed that process monitoring relied upon trends in the engine performance metrics. Data augmentation was performed on all continuous variables to provide trend information to the statistical models. These trend values included measures such as average, slope, deviation, and any shift in these over time. The time period of consideration ranged from prior flight to three years of flights with a number of binning selections in-between.

A variety of model prediction methods and compounded methods were attempted. Data boosting applied for all records related to an event increased the model sensitivity. Data boosting refers to the replication of rare data points to improve the sensitivity of a classification algorithm. Boosting is subject to extreme limitation when data accuracy is in question, and therefore its use was limited to well-known events[45]. The limited number of events on which to train the model required the team to consider many methods of treating the response variable. A variety of decision trees and classification algorithms were attempted to predict probability of failure within a future time period. Alternate approaches used linear, logistic and generalized linear regression to predict days until an event would occur.

Due to the structure of the source data, with very few events available to predict, model specificity or true negative rate was above 0.95 in all testing datasets. As a business objective the model was trained to have a high recall rate, which is the fraction of actual events that are correctly identified as events in a classification model. This value could be controlled to a range of 0.5 to 0.8 in all testing datasets.

2.2.1.3 Model Outcome

The model implementation is ongoing at the time of this thesis. The engines identified as high risk for an event are output by the model with a list of decisions from the tree based model that drove the high propensity rating. This list is reviewed by the engineering and fleet management teams internally before decisions to intervene. This review procedure presents a major transition from reactionary root cause investigation into pro-active maintenance planning. A process for continuous updates is in place that provides new data sources to the model from ongoing event detection in the field.

The false discovery rate on testing data ranging from 0.4 to 0.8 indicates that for every event successfully averted between 1 and 4 investigations were of no direct value. When translated to specific engines, the actual risk can be evaluated by human investigators after relatively minor expended effort. Current estimates for this labor are one full time equivalent (FTE) employee per year for false discoveries, and one FTE for discoveries which result in contacting the operator. Two of the three event prediction models that were funded were discontinued after completion due to poor recall rates on testing data.

2.2.1.4 Failure Analysis

Although the models provided value, they failed to hold up to their promised value. It was possible to foresee, and in some aspects prevent, this problem through better definition of the business value throughout the model development.

Improved data in the form of derived engine health indicators is actively maintained by the company. These data were available at any time throughout the development of the model and were among the sources used for its development, though some data was overlooked due to the difficulty of engaging all parts of teams of the company simultaneously.

During a critical check point on the process funding and model planning, decisions were informed by the use of the model accuracy. As defined, the accuracy of a classification algorithm is equal to true positives plus true negatives divided by total observations. Unfortunately, in rare event detection systems accuracy is not representative of total model performance. For example if only one in one hundred cases is an event, then a pure guessing algorithm that predicts one event incorrectly would have 0 correct events, 98 correct non-events, and 2 incorrect predictions with a total accuracy of 98%. Accuracy was used as a decision making value by default since the business value associated with the false positives and false negatives had not yet been determined. Without that understanding, the balance between precision and sensitivity could not be outlined as a requirement. Better measures of performance before establishing this business value are Informedness and Markedness[46]. Informedness is the true positives divided by total positives, plus true negatives divided by total negatives, minus one. Markedness is the true positives or all positive predictions, plus true negatives over all negative predictions, minus one. While recall in these models had been tuned to 80%, Informedness and Markedness of the original models ranged from 10% to 15%. When this was observed, two of the three models were suspended.

2.2.1.5 Recommendations

Value definition throughout the model development would have reduced costs and improved the outcome of this project. The development of a Data Analytics organization with central responsibility for advertising of data and collection of methods within the company will improve future model performance. An improved integration between engineering staff and data analysis teams would have enabled this type of work to succeed. "You can't do just data mining, it won't make [you] smarter about future situations[47]."

2.2.2 TOW Study

2.2.2.1 Objective Description

The company pursued a Bayes Net regression model for identifying the primary drivers for engine time on-wing (TOW). Bayes Net models enable hierarchical regression where certain factors are given primacy as priors due to system knowledge. The study focused on a fleet of several hundred engines with a mature lifecycle. The objective of the study was to inform future work on TOW studies and expand the knowledge acquired from one mature fleet into a future engine program.

2.2.2.2 Model Development

The development team met on a recurring basis with both engineering and fleet support staff to ensure that all possible data sources were considered and every element of business value from the models was well understood. A large investment was made early in the project in data source identification and rationalization. The team then developed an algorithm capable of forecasting forward the degradation rate of an engine using a Bayes Net. The model derived a prediction variable of the rate of degradation for a key engine health metric. The point at which this metric crossed zero was identified as the expected end of TOW. The model demonstrated an improvement over existing linear models for that engine in two demonstrated test sub-fleets. Existing linear models over-predicted the TOW by as much as 40% and it was for this reason that fleet level linear methods were not in use by the company. The new method of slope prediction corrected for this error and the coefficient of variation from the existing model reduced from 36.8% to 34.5% in one case and 29.5% to 26.2% in the second.

2.2.2.3 Model Outcome

The model was not adopted because the organization could not identify a way to use the model given existing data and processes. Factors known and proven by engineering teams to affect engine health had not appeared as primary variables in the final model. The isolation of variables to those readily measured by the business has not been performed. The airport level severity factors were not applicable to any engines except those on which the model had been trained. As a result, the model relied heavily on the prior knowledge about a fleet performance and lessons from an existing fleet could in no way be translated to a new engine design. The TOW study failed to deliver the necessary flexibility needed to bring value to the business.

2.2.2.4 Failure Analysis

Model flexibility in this instance required that the model isolate the new effects from the existing effects proven by first principles physics. Only in doing this could the model be validated, applied to current predictions, and abstracted to future fleets.

First principles models from Pratt & Whitney hot section durability engineering accurately predict degradation of key engine components as caused by factors such as thrust, flight length and air temperatures. While the basis of the TOW study was to identify new factors, the inclusion of proven first principles should not have been overlooked. If any mission parameter is found to be insignificant in a model it should not be due to omission, but rather due to the identification of a new effect of opposite weight to the physics proven effect. This form of model architecture makes room for the inclusion of new physics to play the role of assumed priors, and builds upon the existing data structure natural to the organization. It is possible that a correlated effect related to airport geometry completely counters the effects of physics on engine damage in some regard. For this new knowledge to provide value to the organization it must be documented with respect to the current baseline so that contracts and pricing changes may be made using a new environmental adjustment. Without this information, the models do not allow for incremental learning and therefore cannot be implemented into an existing business process.

2.2.2.5 Recommendations

Future studies performed must account for known physics factors as forced effects. In this way learning will be made complementary to the existing body of knowledge and not confused with it. Any parameter that is currently a part of the established business process should be regarded as included in the null hypothesis. By construction of a model in this way, the future use of the model leaves room for improved physical understanding of the source data, while providing a clear method for incremental improvement in data acquisition.

2.3 Current Interval Methods

In the current process, TOW estimations are supplemented by active engine health monitoring that continuously reviews the fleet status and identifies risks for early maintenance. This section focuses on the initial prediction made at the point of sale, and fleet level periodic assessments made on existing fleets prior to the renewal of maintenance contracts. The initial estimate uses a combination of engineering and regression models to predict TOW for new engine business (NEB). These models demonstrate reasonable accuracy at a fleet level. To adjust for fleet unique operating conditions that are not fully captured by predictive models, mature fleets with sufficient data are individually fit to a Weibull distribution to supersede the initial prediction. This section explains the current processes in greater detail from the source of the engineering model through to application in both NEB and mature fleets.

2.3.1 Physics Models

For each new engine program, hot section engineering (HSE) performs life prediction analysis using first principles of physics. HSE builds finite element analysis models similar to those described in section 2.1.2. The model uses a set of key variables that define engine operating conditions as inputs: ambient air temperature, flight length, cruise altitude, thrust, airframe, takeoff derate, and climb schedule. Each combination of these variables is referred to as a 'mission point'. The model output for each mission point is a predicted η, and β of a Weibull distribution corresponding to mean failure for the module being analyzed. From a combination of these models for different sections and mission points a lookup table is generated for predicting TOW for an engine model. Throughout the lifecycle of an engine, HSE will continue to monitor and update these lookup tables by adding further physics based failure modes to the source models.

HSE classifies the effect of ambient air temperature independently from the rest of the mission point factors. Air temperature has a slight effect on the inside air temperature of the engine. Commercial jet engines are flat rated to perform equally under a wide range of air temperatures, up to a cut-off point. This means that the engine is capable of higher thrust than is provided in order to compensate automatically for high temperatures, which decrease effective thrust. Above the flat rated point, the engine effective max thrust decreases, thereby reducing the climb speed and extending the time the engine spends performing max thrust climb. The effect that ambient temperature has on elevating internal metal temperatures is relatively low. However, the effect on climb time extension is specific and predictable. Longer climb at max thrust is equivalent to running the same mission point simply for a longer period each flight. For these reasons, temperature effects are handled independently of the rest of the models and a temperature based life factor is generated independent of the other factors in the mission point.

2.3.2 Table Application and Severity

Engine Services (ES) receives the new tables from engineering for each engine model and performs model verification and severity calculation prior to applying the new tables. ES gathers fleet operating data for the existing engines when possible. This information is grouped by sub-fleets binned by both operator and airframe and reviewed for shifts in trends over several years. If sufficient data is available for a sub-fleet, ES computes an average thrust, flight length, temperature and cruise altitude for the fleet based on historical data. These average values are checked against the HSE tables to identify the expected TOW for the fleet and this expected value is compared with historical observations. Historical observations for time on-wing are categorized by removal reason to identify data censoring and truncation. Next a Weibull distribution fit provides the method of calculating fleet level average interval. It is generally observed that the engines are either removed near engineering prediction or significantly early due to unexpected causes; the result is an average overestimation of life expectancy and the ratio of actual life over expected life is referred to as the 'severity'. A senior technical fellow in the engineering department states that "Every field problem, every removal is due to variability[47]." ES documents the residual error between the tables and historical observations.

For the purpose of this thesis, the data are categorized into regions defined as continental or sub-continental areas and airlines are grouped by nation of origin. This enables the maintenance policies and general climate effects of an area to be loosely accounted for. A regional correction table generated from these data points forms the basis of new models used for all engine TOW assessments. For each region one adjustment factor is determined which minimizes the average error within the region. The development of the regional severity factors corrects for a number of missing variables in the physics based analysis in a way that is operator independent, while adjusting broadly for political and environmental drivers of performance. Typical results for a selected group of fleets shown in Figure 4 demonstrate an improvement in total forecast accuracy achieved by implementing the regional adjustment. The regional correction effects are shown in Figure 5.

Figure 4: Regional severity adjustment to standard engineering model with 95% CI of Weibull mean

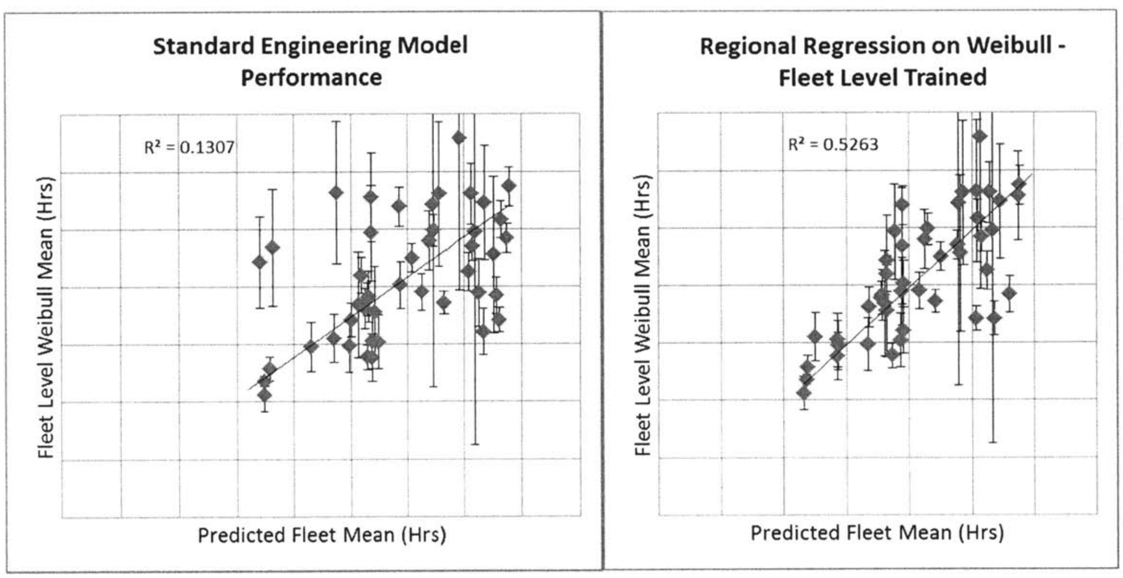

Figure 5: Effective B50 life factor of regions within comparable engine models

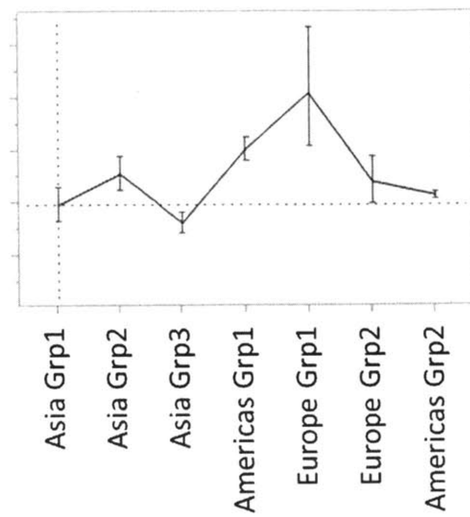

The use of average fleet level values in deriving the severity factor constitutes a loss of information, and a misuse of the HSE tables. The HSE tables exhibit clear non-linearity in their outcomes as a function of each mission point variable. For example, flight lengths of 1 hour, 2 hours and 3 hours may result in a predicted TOW of 5,000, 10,000 and 12,000 hours respectively. Flying a mission in which 50% of the flights occur at 1 hr length, and 50% at 3 hour length would yield a final interval of 8,500 hrs if treated independently, or 10,000 hrs if the flight lengths were first averaged and then applied to the HSE tables. Despite this, the combined model exhibits better overall data fitting due to its ability to over-fit the regional severity in a way that corrects for its statistical impurity.

2.3.3 Weibull methods for existing fleets

The application of the engineering tables and severity adjustments forms the basis of understanding engine life prediction prior to the first shop visit for any new fleet. A periodic review of fleet performance against this expectation ensures that company is able to properly schedule maintenance work, update regional severity correction factors, and deploy corrective action for damage that arises faster than expected. When full intervals have been observed for at least three engines of similar model and thrust, a Weibull analysis is performed on the set of all engines in the sub-fleet of engines with the same operator, configuration and thrust rating. A review of the shop visit notes ensures that the removals are driven by engine deterioration identified during inspection rather than "events" such as foreign objects, bird strike or service bulletin incorporation, collectively classified as censored observations in Weibull analysis. Removals due to life limited part replacement schedules represent truncation in observable time length and are also classified as censored records. The Weibull mean and

95% confidence interval (CI) are compared to the engineering prediction. When the Weibull mean CI is outside the value predicted by regional severity and engineering tables the fleet forecast is updated to reflect the new Weibull mean rather than the predicted value. This analysis for a single fleet can be done in under an hour after the data are collected.

When insufficient observations of non-censored engine refurbishments have occurred, Weibayes methods provide validation or challenge to the baseline fleet prediction. Weibayes is a statistical method involving data where only censored values are available. Either the shape or scale parameter of the Weibull distribution is assumed. Next Bayesian inference is used to identify the lowest possible value of the unconstrained parameter that permits a 95% probability of all known events simultaneously not occurring. ES performs Weibayes using the same off the shelf software that performs the Weibull analysis. To provide a conservative estimate of life a Weibull beta typical of tightly grouped failures is assumed for the Weibayes model and the 95% confident lower bound of the mean is compared to the original expected fleet performance. The Weibayes derived fleet level lower bound mean is generally trusted when it exceeds the originally predicted value. When the total fleet has a very low number of flight hours, typically less than half of the predicted flight hours, the Weibayes analysis generally provides a mean confidence interval too broad and a lower end too low for general acceptance. Personal judgment by the ES analysis group in collaboration with the Pratt & Whitney customer fleet manager results in a decision to either use the Weibayes value or continue to base forecasts on the engineering tables and regional severity.

The combination of engineering tables and severity factors provides value to the company in estimating cost structure for both new contracts while the adjustments using Weibull and Weibayes analysis prevent or minimize the financial effect of initial forecast error. At the same time the analysis incurs minimal cost for either data collection or analysis as the data is already recorded and analysis can be performed within only a few hours per contract. The model is fully flexible to the needs of the company because it does not rely upon special knowledge. The engineering tables and severity grids are applied multiplicatively as independent effects. The Weibull analysis within off the shelf software can be done in a repeatable and easily validated way after minimal training or statistical knowledge and is used as the basis of estimate for all fleets in place of the grid prediction once sufficient data has been observed. The application of the Weibull analysis and tables are flexible to unique aspects of the fleet management strategy or rapid changes in business knowledge. Overall the model system provides strong value and any improvement in model accuracy must be accomplished with similar ease of use in order to be accepted by the company.

3.0 Model Development

A three layer model is developed to describe the effects of (1) engine physics, (2) inspection and maintenance policies, and (3) environmental effects. The model pre-development process relies heavily upon engineering informed system dynamics. This review ensures that the data systems used in the model are complete and sufficient for describing the problem and identifies early concerns in the data structure that inform later modeling decisions. The physics of underlying systems, combined with input from the system dynamics view, contribute to the development of new variables from the source data and the selection of proper units for comparison between data points. Censoring and truncation present in the problem are described and multiple approaches for analyzing data within this context are proposed. The individual performance of each model layer is analyzed. The final model is abstracted to the airport level to develop a visualization of the aggregated effects on a global scale. Finally, the sensitivity of the system to certain statistical methods is reviewed. The model is found to significantly out-perform current methods. Changes to the treatment of leverage points are found to be significant in enabling model convergence.

3.1 Overview of approach

Engines experience a wide range of operating conditions due to effects of seasons, operator policy, airport topography, and many other sources. Such broad datasets present a challenge to model development. Many conditions cause different types of overall degradation to the engine. As a general rule, degradation by abrasion, fowling, misalignment, or bleed valve deterioration cause an increase in energy loss along the direction of thrust. The engine computer compensates for this energy loss by increasing fuel flow to achieve the target engine pressure ratio. Increased fuel flow results in higher operating temperatures, resulting in higher degradation to hot section components. For example turbine blade creep can be effectively modeled using time at temperature[86]. Therefore, increased temperature during operations reduces the overall time to creep failure. Oxidation, sulfidation, and metal fatigue all increase temperature, and lead to further reduced gas path efficiency. For this reason, the airline industry commonly uses exhaust gas temperature (EGT) as a bellwether for the overall engine health[11].

Although not all engines are removed due to an EGT degradation, the majority of degradation modes lead to increases in EGT during normal operation. Therefore a system dynamics model built around the EGT provides insight to the sources of damage. System dynamics is applied in three layers. First, every potential effect on the engine is added to the diagram to determine the breadth of the problem statement. Second, a process flow of available data is developed and the implications of these data origins are discussed. Finally the process flow is used to inform pruning of the first model and highlight areas for future data generation. The resulting model represents an important aspect of data rationalization that is critical to developing causal chains and preventing over fitting in the final model.

3.2 System Dynamics Model

The following system dynamics models were developed through a series of interviews with engineers regarding failure drivers in the engines. A number of effects that appear in the model are not directly quantifiable. Once all factors have been identified, we rationalize the results through an iterative process of confirmatory tests. These tests support the model in each area where data can be examined to validate directionality. In certain areas aggregation of factors reduces complexity while retaining causal information. This is prevalent where driving factors are overlapping and underlying causes are well understood.

To develop the system dynamics chart shown in Figure 6 performance characteristics from all modules of the engine were used combined with environmental factors throughout a flight. Flight phases considered for effects were taxi, takeoff, climb, cruise, landing, and storage. The model is constructed around the EGT as a proxy for overall damage to the engine. Environmental factors fall into two groups: parameters associated with the air intake; parameters from flight mission. Human factors are applied to the model based upon interviews with the Pratt & Whitney maintenance cost group. The model considers the effects of a single flight through all phases. The objective of the model is to inform the factor selection in predicting the time between shop visits. Therefore maintenance policies of the airline are also included. Although they do not contribute to damage, they do affect the detection of this damage. In certain cases airline policy may lead to over inspection and identify minor damage before it has led to significant degradation. If continuous and targeted maintenance action is taken, it will reduce overall damage aggregation and may affect the time between overhauls.

Figure 6: System Dynamics view of engine effects during takeoff

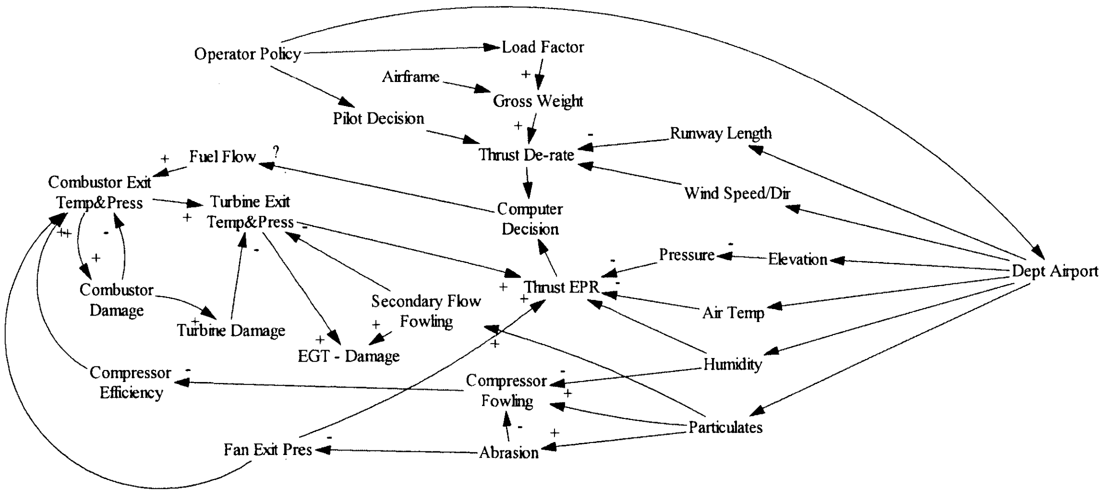

Pratt & Whitney acquires engine performance data through multiple paths, and data conditioning is applied at different steps depending upon the origins of the data. ACARS messages transmit engine performance data to ground stations in a fixed code four characters per measurement. ACARS messages send data from multiple flight phases when stable flight conditions exist for that phase. In some cases where stable conditions cannot be achieved, data points may be skipped or the report will send limited information. Messages are addressed to multiple parties which may include the engine manufacturer depending upon the contract. When the data is addressed to Pratt & Whitney, the company collects the data to a central database directly from the ground stations immediately. When the data is not addressed to Pratt & Whitney, the airline collects the data and performs internal curation. This curation may involve data compression, unit conversions, or similar changes to the data which retain the overall information in the data. For example, certain values transmitted by the engine are considered redundant when operating properly. Therefore the airline may choose to store only one of these values after confirming normal operations. On pre-agreed intervals, these airlines transfer all collected data to Pratt & Whitney which uploads the data into the same central location as the data from airlines that addressed Pratt & Whitney directly. While data within a single airline is consistent in formatting and update frequency, data between airlines may be inconsistent until corrections are applied for different units and data compression applied by the airline. Both the availability and cross fleet consistency of data informed the down selection of variables for the final screening model.

As a final consideration in developing the model architecture, we consider the variables that can be known before the flight and avoid the use of in-flight variables such as vibration levels. For the purpose of the severity study of interest to Pratt & Whitney, these parameters need to be known years prior to the flight. These include airport parameters, hours to cycle ratio, planned cruise altitude, operator maintenance policy, engine configuration and thrust selection which are roughly designated by the maintenance contract. The system dynamics flow demonstrates that in each flight phase these values, which are highlighted with boxes, are well positioned as input drivers to the entire model. Although other variables may improve the understanding of engine damage, these values cannot be well known at the time of contract sale. Future models are discussed in later chapters which may include these variables.

Figure 7: System Dynamic view with primary knowable factors highlighted.

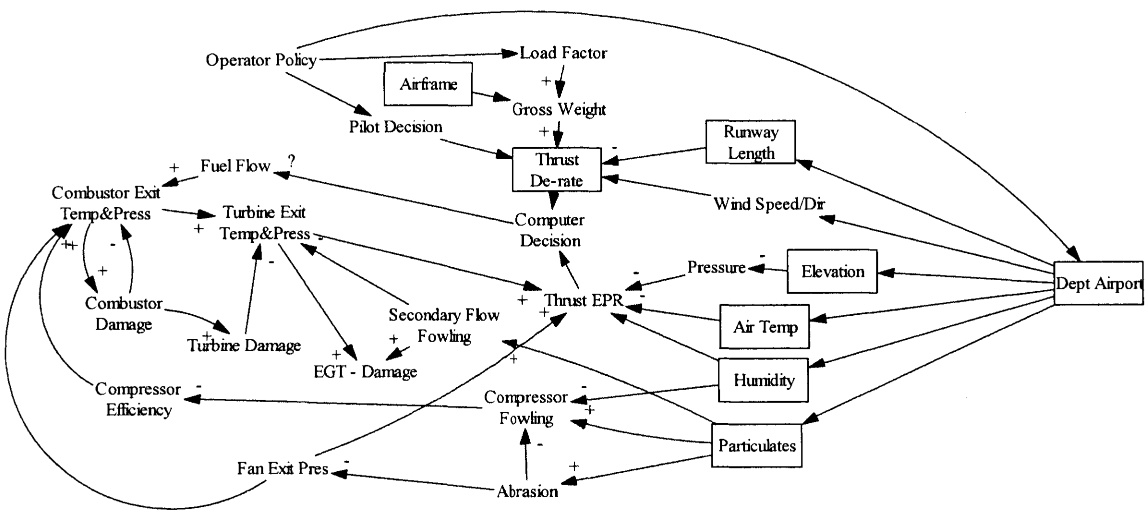

3.3 Physics interactions and Data Augmentation Applied

The down selection of variables performed in prior section loses a great deal of information regarding the internal physical interactions of the engine. First principles models that describe the internal workings of the engine based upon these key inputs restore this information to the model. The Pratt & Whitney Hot Section Durability team performs a variety of 1-D and 3-D simulations that identify the relationship between hot section damage and the typical operating characteristics of the engine. Failure is defined as the exceedance of the engineering defined limit in any of these degradation modes. This engineering defined limit is used as the guide for development of maintenance manuals and the simulation model informs the development of the inspection plan for each engine module.

As inputs, the model uses environmental factors and operating parameters together referred to as mission parameters. A single set of mission parameters is one "mission point". For each mission point expected to be observed in operation, the simulation is run and the mean failure point is recorded. The resulting table provides a relationship between common mission points and the expected engine time on-wing collectively called "HPT Grids". The simulation does not provide a continuous dataset due to the computational time required for each mission point. The model put forward in this paper takes these HPT Grids and performs multi-dimensional linear interpolation to produce a time on-wing estimate for every possible mission point contained within the limits explored by the HPT Grids. In addition, the HPT Grids are extrapolated linearly along a single axis for the cases where one parameter of a mission point observed in the fleet exceeds the range simulated in the HPT Grids and the other parameters of the mission point are within the simulated range. In approximately 5% of missions observed, more than one parameter was outside the simulated range and the mission is disregarded.

Through this method every observed flight is related to the HPT Grids which provide an expected number of flight hours until failure at that mission point. The ratio of the number of flight hours in the flight to the expected flight hours to failure from the HPT Grids yields an approximate percentage of life used by the flight which may be added cumulatively through an engine's life. Where flight data is missing due to filtering or data loss, the average life consumed per flight by that engine in all observed flights is used instead. The cumulative life used by the engine as defined by this method is referred to as the Life Factor for that engine and is reset to zero when the engine undergoes an engine level overhaul.

In considering methods of grouping data for TOW analysis, the prediction variable and sampling bias are taken into account. The analysis of engine damage drivers indicates that the majority of failures are caused by cycle driven damage rather than flight hours driven damage. For this reason the prediction variable designated for the majority of models is the cycle count at failure. One aggregation method is to consider all values having transpired over a single interval as a single event. This model is useful in considering how the total history of the fleet has behaved in a comparative interval to interval way. However, this method averages values across an uneven basis. Engines with more flight experience before removal have a greater number of points aggregated. This decreases the effect that a leverage point within that flight history could have in any aggregation method. This limitation is countered by the use of cycle weighted standard error, and is discussed further when considering sensitivity of the model.

The chosen method of data aggregation depends upon the type of data being analyzed. Aggregation of the average value across a set of flights may eliminate valuable information about the spread of such values. Simple expansion to the aggregation of higher moments is not sufficient for some applications. When the expected response between the parameter and the target variable is linear, average value may be safely assumed. When the responses are non-linear, averages may only be taken after the non-linear transformation has occurred. Higher order terms are created only when the physics of the behavior measured suggests that the product will have physical meaning. Otherwise the decision of how to aggregate each field is made after effect screening has occurred. The individual factor residual plot is analyzed to determine if the response appears linear or non-linear. Before non-linearity can be assumed the correlation matrix for the factor of interest is examined due to the complexity of the system being modeled. If two variables exhibit non-linear residuals, cross terms of those variables are tested.

3.3.1 Shop Visit Data

Shop visits are recorded by Pratt & Whitney on the engine using work scopes from 0 to 3 indicating the amount of work performed on each module of an engine where 0 indicates no work, 2 indicates major repair or replacement, and 3 indicates a full overhaul of the module. For the purpose of this study an engine level overhaul is defined by a work scope of 3.0 on any key module or work scopes greater than or equal to 2 on to all four key modules where the four key modules are the Combustor, High Pressure Turbine, 1st Nozzle Guide Vane, and Low Pressure Turbine. This definition is borrowed from internal convention at Pratt & Whitney. Each period of flights separated by an overhaul is referred to as one "Interval" for which the time on-wing is defined in either flight hours or flight cycles.

3.3.2 NASA Data

Data is acquired from the NASA TERRA satellite and converted into monthly average volume mass loading over each airport in the world each of eight aerosol types shown in Table 2. Details regarding the data acquisition and conversion are provided in the Appendix. Since the overpass of the satellite is performed at approximately solar noon in each airport, there is no ability to correlate the data to the exact flight takeoff time. The data is used to analyze seasonal changes in regional aerosols over the takeoff locations. The final dataset by airport and orbit number is maintained in a Teradata 14 database with all TERRA records from February 2000 through June 2014.

Table 2: Aerosol types modeled by NASA MISR

Aerosol Code	Type	Size (μm)	Expected Composition
1	nonabsorbing	0.06	sulfate/organic
2	nonabsorbing	0.12	sulfate/organic
3	nonabsorbing	0.26	sulfate/organic
6	nonabsorbing	2.8	salt/organic
8	absorbing	0.12	sulfate/organic (ssagreen.9)
14	absorbing	0.12	sulfate/organic (ssagreen.8)
19	grains	0.75	mode1 dust
21	spheroidal	2.4	mode2 dust

3.3.3 Exhaust Gas Temperature

The EGT may have a hard limit beyond which the internal metal temperatures are considered unsafe. Engines which reach this limit will be removed for overhaul. In order to ensure that this number is never reached, the engine computer calculates a realistic worst case (RWC) EGT that could occur under the flight conditions present during takeoff. This RWCEGT is adjusted in units to standard day and atmosphere. In practice, individual values for RWCEGT may exceed the engineering defined limit. When several such takeoffs occur, the engine is removed for maintenance. In some engine families the EGT limit is the primary reason for maintenance intervention, while in other engine families the EGT limit is not reached before other parts fail inspection. In both cases EGT increases are monitored as one indicator of health. The difference between the RWCEGT and the EGT limit is plotted in degrees Celsius as a positive value approaching zero as the engine degrades.

3.3.4 Air Temperature

Ambient air temperature is corrected to International Standard Atmosphere (ISA) by the following equation in order to account for the effects of temperature and pressure jointly:

Equation 2 $$\Delta \boldsymbol{T}_{\boldsymbol{amb}} = \boldsymbol{T}_{\boldsymbol{static}}(°\mathrm{C}) - (\mathbf{15} - \mathbf{0.0019812} * \boldsymbol{Alt(ft)})$$

where the result ΔT_{amb} is referred to as the delta between ambient temperature and ISA (DTamb), and enables a conversion from any temperature and pressure to the pressure equivalent temperature at sea level.

3.3.5 Data Aggregation

The volume mass loading of particles over each airport is observed to follow a log normal distribution. To support generalized regression studies, each parameter in the atmospheric dataset was converted from absolute units of µm into lognormal quantiles. Where the value observed was zero, a quantile of -3 was assigned as filler. The aerosol data was further binned into 5 equal width groups on the µm where the high end of the bin was the intercept of +3 quantile to the µm scale. Industry established prediction equations for the lifetime of an engine are based upon flight length, takeoff thrust, and ambient temperature at takeoff. In addition to these, the cruise profile affects the temperature and thrust assumptions during climb which are required for damage simulation. Simulations by Pratt & Whitney engineering are performed on a single mission basis. The mean time to failure is predicted by assuming a fixed set of parameters for flight length, takeoff thrust, temperature and cruise profile. Common physical failure modes of the combustor and high pressure turbine are included in the model. To apply these predictions to engine performance data, the predicted interval is established for each takeoff and the average of predictions is identified for each full maintenance interval. A full maintenance interval is defined by shop work scope performed as either a complete replacement of the combustor, or a major overhaul of all hot section engine modules including the high pressure compressor, combustor, nozzle guide vanes, and high pressure turbine. All other attributes recorded by flight are aggregated over the engine interval and collected into average, standard deviation, average quantile, and percent of flights in each of 5 equal width bins.

3.3.6 Airport Severity Factor

The thesis objective is to develop an airport severity factor (ASF) to work in complement with engineering prediction values. In this way the ASF will be most transferrable to new engine programs and contract language. For each engine interval observed, a life factor is computed by dividing the actual observed flight hours by the physics derived mean time to failure. The resulting unitless value ranges from 0 to 1.4 shown in Figure 8. The engineering predictions generally establish an upper limit to engine lifetime. With respect to the engineering predictions, engines can be seen in four categories: those that deteriorate early; those that have not yet deteriorated; those that are removed without deterioration; those that deteriorate due to modelled causes. The fourth category imposes an upper limit on life factor, while the other three groups tend to occur before the engineering predicted limit. Engines that have not failed are coded as right censored observations along with engines that are removed with life factors below 1.0 that have not deteriorated. Engines that are removed prior to deterioration with life factor above 1.0 are considered data truncation, generally occurring due to life limited part replacement schedules.

Figure 8: Life Factor distribution of all deterioration caused hot section refurbishments

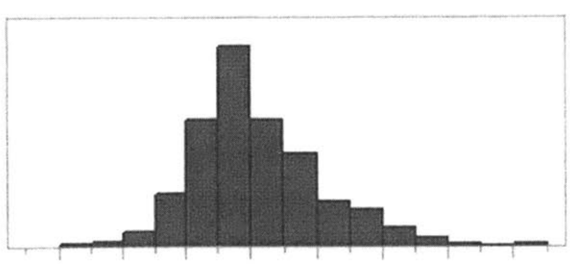

3.4 Primary Effects and Screening Methods

Configuration of the engine provides the largest single effect on engine lifetime with variance between configurations in excess of 50%. Four differentiating characteristics are identified for the engine. Three of these configuration differences in the fleet arise through maintenance intervention throughout the lifecycle as new technologies are implemented to the fleet, while the fourth differentiating characteristic is the engine age. These upgrades are placed on new fleets, during overhaul, or as needed to accommodate for harsh environments. Data observation sizes are shown for each configuration combination in Table 3. The engine age since new, after controlling for configuration can be measured in multiple ways: cycles, hours, shop visit number, or mature run vs. first run. All tests for effect are done by both considering the effect on cycles between shop visit and the effect on life factor after mission adjustment. Since the life factor mission adjustments are developed from first principles physical models, adjusting for them before other analysis gives primacy to a known interaction and is not considered to be a loss or confusion of information.

Table 3: Fleet Configuration Counts in Available Data

Configuration	Full Observation	Censored or Truncated
First-FFF	17.6%	3.5%
First-FTF	5.6%	10.0%
First-FTT	0.0%	0.2%
First-TFF	0.0%	0.4%
First-TTF	0.4%	19.3%
First-TTT	0.0%	9.1%
Mature-FFF	8.4%	14.9%
Mature-FFT	0.0%	0.1%
Mature-FTF	1.5%	20.0%
Mature-FTT	0.0%	0.7%
Mature-TFF	0.0%	0.1%
Mature-TTF	0.0%	5.3%

Configurations with insufficient observations are removed from study. For the purpose of fleet simulation, individual studies of these configurations are performed at the time of technology deployment. These studies yield a predicted life extension effect, and are used to apply all models from this study to those fleets outside the technology range studied. A survivability analysis is performed on each of the remaining configurations both before and after adjusting for engineering predicted life.

Three distribution types were tested for application to the engine survival modeling: Lognormal, Loglogistic, and Weibull. Each of these models is formed from a different assumption about the underlying failure mode. Lognormal distributions arise from the convolution of the exponential distribution with the normal distribution. Problems fit well by lognormal distributions are systems with an underlying exponential component, essentially random arrival in a time based process, with a normal distribution on quantity or duration or vis-a-versa. For example, aerosol mass loading as a function of particle size is lognormally distributed since the time duration of an aerosol can be described with an exponential decay and the elevation of these particles into the atmosphere is caused by an essentially normal random process in steady state[48]. Loglogistic distributions are used to define survival models with an accelerating mortality rate before the peak and a declining mortality rate after the peak[49]. The behavior is similar to the lognormal distribution except that the underlying logistic distribution has heavier tail contribution than the normal distribution. The loglogistic distribution will fit better than lognormal in conditions where extreme values are being reinforced away from the mean. The Weibull distribution is commonly used in engineering survival analysis because it describes the sum of minima's of a distributed event[50]. This condition applies to survival analysis where the underlying failure mode is assumed to be normal, and since the nature of the failure ends the experiment, the resulting distribution of repeated trials is the sum of minimum occurrences.

The life distributions based upon Lognormal, Loglogistic, and Weibull are parameterized across the configuration factor to assign a value to the configuration effect. For the time to event, the analysis is repeated using both the number of flight cycles between overhaul, and life factor. Life factor refers to the ratio of cycles to expected cycles based upon engineering prediction for that interval. The B50, or median failure point, of each configuration is plotted in Figure 9 and the ratio of these values is shown in Table 4. Table 2 displays the values from the Loglogistic survival analysis which was the method of best fit with AIC 618 compared to Weibull and lognormal with AIC 738 and 795 respectively. The effect of the configuration upgrades is seen to depend upon the way flight mission is handled. For example, the First-FTF configuration flies more cycles than First-FFF, and experiences longer life factors. This reflects an underlying relationship between configuration types and variance in flight mission and therefore life factor is a preferred method of analyzing configuration differences. It is not obvious that the mature run on FTF should be longer than the first run of FTF while mature FFF run is shorter than first run FFF. If the effects of mature run are believed to be the result of non-critical part degradation, and the difference between FTF and FFF affects only critical hot section parts, then the effects of mature run would be expected to be similar for both configurations.

Figure 9: B50 Life with 95% confidence for each configuration as calculated by three distribution methods

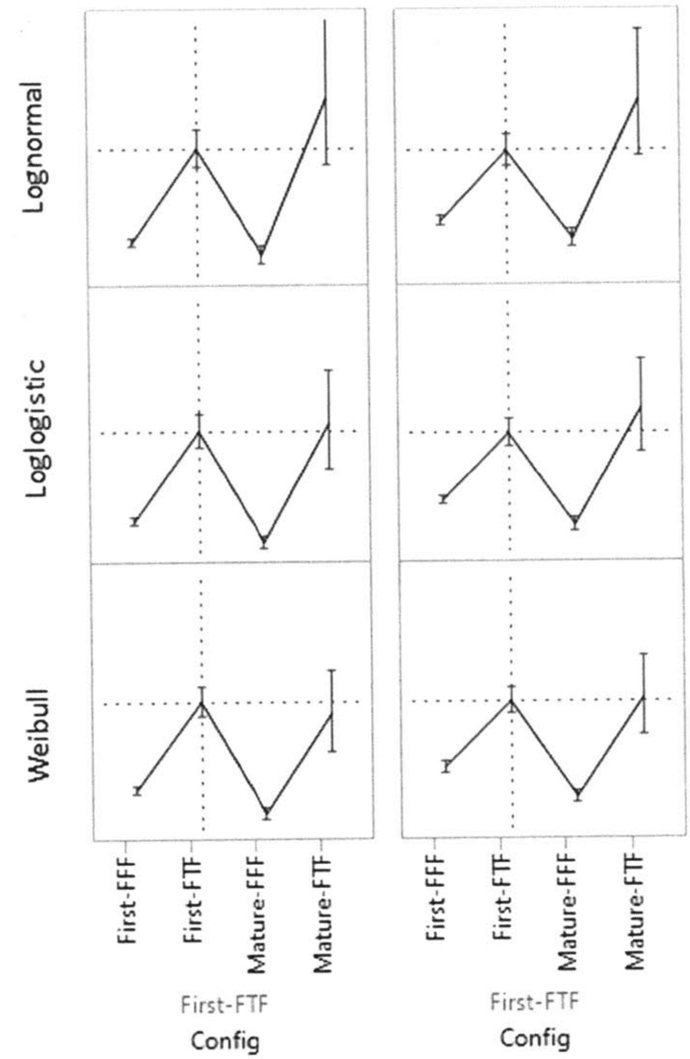

Table 4: Ratio of B50 life between configurations for both cycle based and life factor based derivations

		First-FFF	First-FTF	Mature-FFF	Mature-FTF
Cycles, Loglogistic B50 Ratio	First-FFF	1.00	0.67	1.14	0.65
	First-FTF	1.50	1.00	1.71	0.98
	Mature-FFF	0.88	0.59	1.00	0.57
	Mature-FTF	1.54	1.02	1.75	1.00
Life Factor,	First-FFF	1.00	0.74	1.16	0.67
	First-FTF	1.36	1.00	1.57	0.91
	Mature-FFF	0.86	0.64	1.00	0.58

Loglogistic B50 Ratio	Mature-FTF	1.49	1.10	1.72	1.00

The effect of engine life on future interval can be carefully deconstructed by plotting a parametric survival analysis on several definitions of engine life. After plotting survival analysis to lognormal, loglogistic, and Weibull distributions parameterized by interval number, a ChiSquare test of regression indicates a probability <0.001 of a distinction between single variable regression in the location parameter and separate location for each value of interval number. The higher interval numbers have intervals exceeding that predicted by regression requiring each interval to be treated independently for best fit. This occurs in lognomal, loglogistic, and Weibull survival distribution fitting. The engine age at which the interval length increases for FTF configuration is at interval number 3, while FFF does not increase until interval number 4. One possible reason for this difference is that the mean interval of FTF is longer than FFF, and when the engine age used for comparison is cycles or hours, there is no statistical difference between the age where FFF and FTF exhibit increasing interval. Combined with an understanding of the nature of physical differences between FFF and FTF this indicates that the observed behavior is unrelated to the configuration distinction and is an age related behavior alone.

Engines contain parts with pre-determined cycle limits before they must be replaced which are collectively called life limited parts (LLPs). When the engine overhaul is performed near the end of LLP life, the LLPs are proactively replaced to prevent them from causing another overhaul too early. For example, if the engine is inducted to the shop for its first shop visit with 2,000 cycles remaining on the LLPs and the engine is expected to fly for at least 12,000 cycles then the LLPs will be replaced early. When the LLPs are replaced, the engine undergoes a more thorough overhaul than when LLPs are left in place. The difference in interval can be seen by plotting the mean achieved interval as a function of LLP life used at the beginning of that interval. Generally LLPs are scheduled for removal every other overhaul. When the percent of LLP life used at overhaul is greater than 50% then the LLPs are more likely to be replaced, while if the LLP life used is less than 50% the LLPs will not be replaced. Figure 10 shows that the relationship between LLP life used and mean achieved interval is generally linearly increasing after passing through a minimum at 25%. We consider overhauls as major if LLPs are replaced and minor when LLPs are not replaced. Table 5 shows that this consideration of major and minor overhaul shifts succeeds in decoupling the effect of engine age from engine configuration. Both a major-FTF and minor-FTF have 41%-42% greater interval than the respective FFF build and major overhaul within the same configuration exhibits a 19%-20% improvement over minor overhaul regardless of configuration. This decoupled relationship is not only easier to understand and use for predictive work, it is also more likely to be an accurate representation of the real world. Despite the fact that both this relationship and the system portrayed in Table 4 are of equal statistical strength, the figures in Table 5 are a more compelling explanation of the data and are used for configuration adjustments in the remainder of this paper. Figure 11 demonstrates the relationship between the four categories established and the engineering predicted limits. This chart includes only those engines that were removed due to deterioration, and does not account for censoring or truncation, so it may only be used to demonstrate that the engineering prediction behaves as a proper upper limit to all four categories. The comparison between Figure 10 and Figure 11 highlights the importance of including the censored and truncated records in the analysis in that simple regression methods looking only at the failures would indicate a similar behavior of Major_FTF and Major_FFF while survival analysis accounting for censoring yields a significant distinction between the two configurations in Figure 10.

One additional factor contributing to the overall life similar to configuration and build plan is the maintenance work occurring between overhauls. We compare the B50 point of fleet life factor by amount of total work performed. This analysis can also be applied to the work performed by section within the engine for further insight to the module level effects on overall performance. However, general application of this information is limited by the high collinearity between maintenance performed in nearby modules. The resulting data in Figure 12 provides value in both understanding biases in the data before further analysis and in establishing the return on investment for maintenance actions.

Figure 10: Achieved interval duration compared to engine age at the start of the interval.

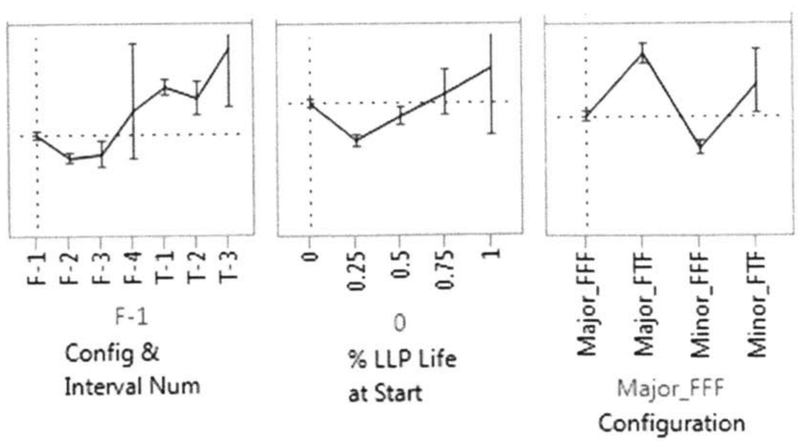

Table 5: Ratio of B50 life between configurations with major or minor overhaul

		Major-FFF	Major-FTF	Minor-FFF	Minor-FTF
Life Factor, loglogistic	Major-FFF	1.00	0.70	1.19	0.84
	Major-FTF	1.42	1.00	1.70	1.20
	Minor-FFF	0.84	0.59	1.00	0.71
	Minor-FTF	1.18	0.83	1.41	1.00

Figure 11: Engine Interval vs. engineering model prediction by starting conditions and configuration

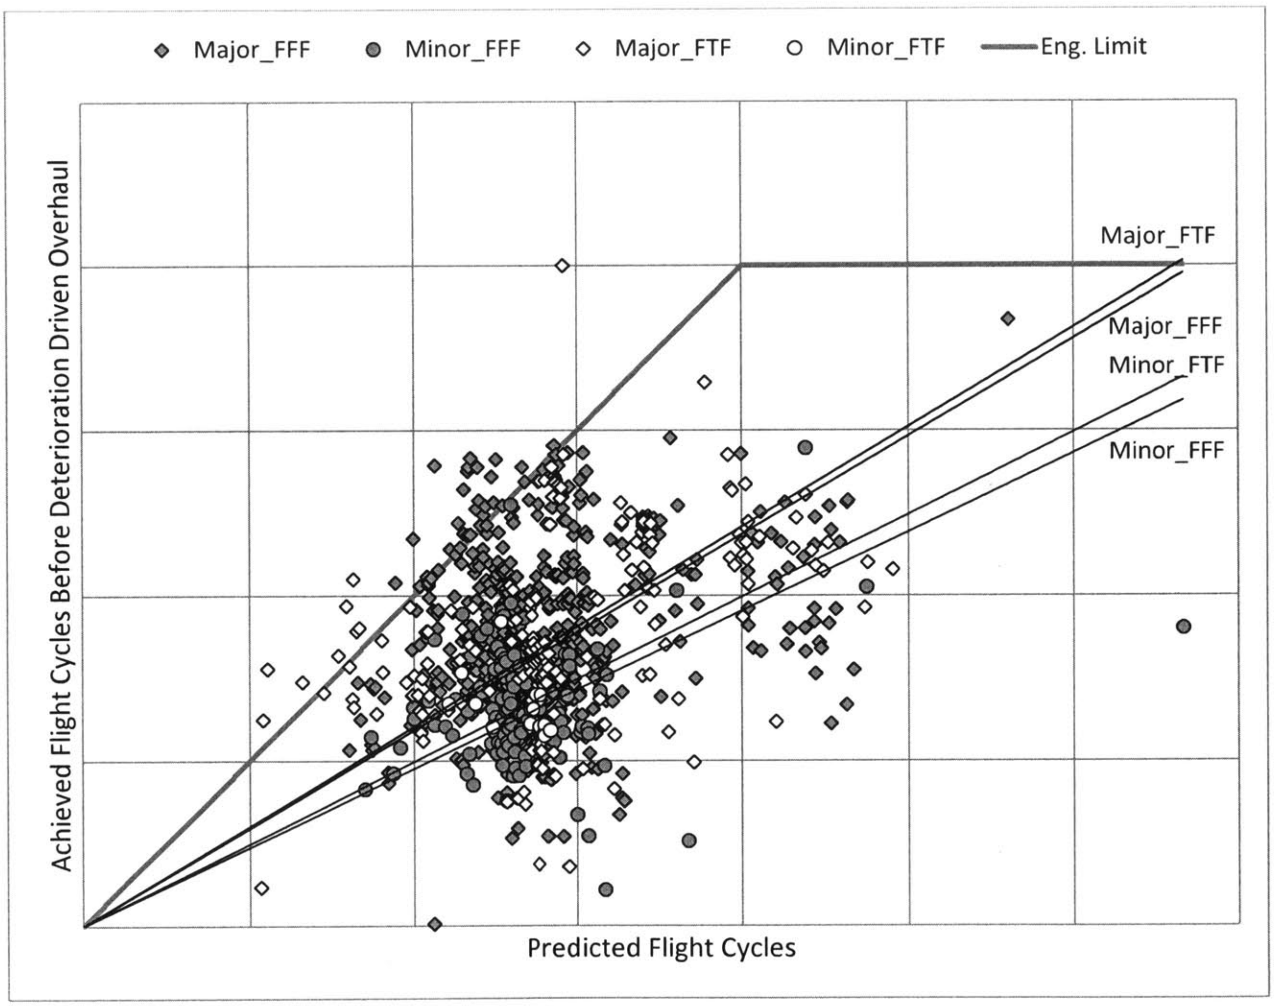

Figure 12: Relationship between B50 Life Factor achieved and level of maintenance work

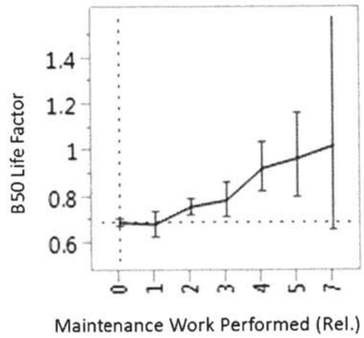

The model hierarchy must first compare data on an equal basis in as much as the basis of engine performance can be known a priori. This allows for a study of atmospheric effects to provide transferrable knowledge to future engines, and to provide insight with regard to the original basis conjectures. The life factor of each engine is adjusted using the factors in Table 5 which enables testing of further data on the collected dataset. In doing so we rely upon the product knowledge of Pratt & Whitney engineering that these four configurations described should not respond to environmental factors in any materially different way. This assumption underpins the collective use of the adjusted data. Figure 13 shows the data on a loglogistic survival plot both before and after the location adjustment. Before adjustment, distinct lines are present, while after adjustment the distributions are visually similar and loglogistic fits without parameters have lower BIC (550) than parameterized versions with either location alone (570) or location and scale (566). This provides limited affirmation of the assumption that the four configurations after adjustment may be treated as a common configuration.

Figure 13: Loglogistic survival plot of engine life factor by configuration and starting condition before (left) and after (right) adjustment for starting factor based upon loglogistic B50 of each population.

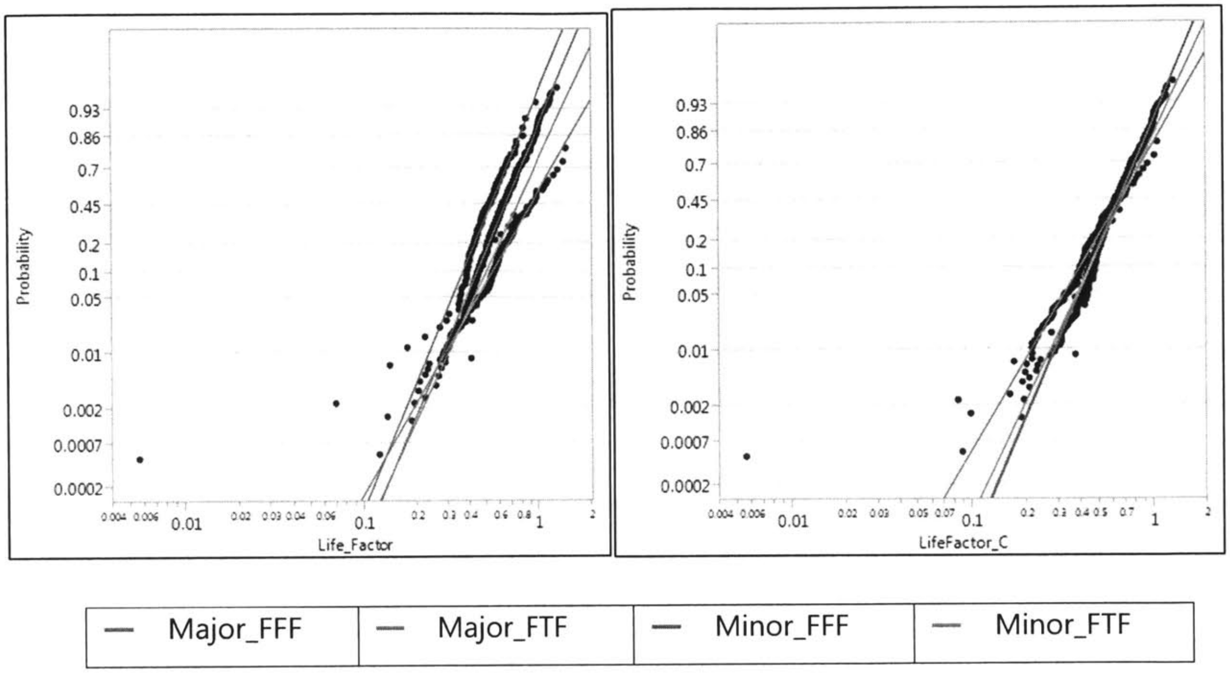

We use the life factor to test the overall performance of the engineering model separately for each configuration group. The model is constructed using thrust, flight length, and ambient temperature. Although climb profiles are included in the model, the optimal profile is assumed. The loglogistic B50 by life factor is measured across each of the three parameters used in the engineering analysis independently on all four configuration sub-groups. The null hypothesis states that the engineering provided life estimates are correct. The null hypothesis is rejected if the B50 of the adjusted life factor is substantially different as a function of the parameters individually or as a whole. When the B50 is greater than average for that sub-group the engineering model is underestimating the life, and when the B50 is smaller, the model is overestimating the life. Figure 14 shows the results of the tests performed. The results across flight length for Minor_FFF appear to be a linear trend in which life expectations are underestimated at low flight hours and overestimated at high flight hours. However, this trend is not repeated in the other configurations and it is possible to draw a slope zero line through all 5 confidence intervals. The slope on major_FTF is more convincing, albeit opposing the slope on both FFF configurations. The 3 hour flight length test point for major_FTF contradicts the trend and indicates that the observed slope may be caused by other factors. The thrust percent derate is based upon the average percent below maximum thrust at takeoff. This is calculated on a fleet level using a maximum thrust for the engine model. The models support engineering predictions regarding thrust percent derate in that these factors not survive null hypothesis testing of the stepwise regression. The temperature effect on engine life demonstrates a trend throughout all four models of elevated values at 10 and 15 degrees above ISA. Temperature effects are covariate with a number of engine damage drivers identified in the system dynamics model in Figure 6. A correlation matrix of primary variables used in the analysis is included in the appendix. Temperature is directly related to humidity and aerosol density. These are collinear to airport elevation insofar as high elevation airports tend to be hot. This is shown in Figure 15. The relationship between temperature and elevation is not universally true of all airports, but is true within the sample set of airports covered by engines in the period of study. Airport elevation in turn causes higher takeoff thrust due to thinner air, and has lower climb time which may be associated to engine damage. The directionality of the effect of elevation alone is difficult to estimate since the decrease in climb damage may be counteracted by an increase in takeoff damage. A parametric loglogistic fit of life factor to average takeoff elevation demonstrates a statistically significant rise in expected life as a function of elevation shown by Figure 16. This may be due to an error in the temperature model, the true effect of decreased climb time, or other effects as yet unknown. First principle simulation of engine parameters was performed to create the temperature response curves will be required in order to isolate the variable effects further. After considering effects that may impair our ability to judge the temperature aspect of the engineering model, there is not sufficient evidence to reject the engineering based predictions from the data.

Figure 14: Test results of engineering first principles models by configuration type

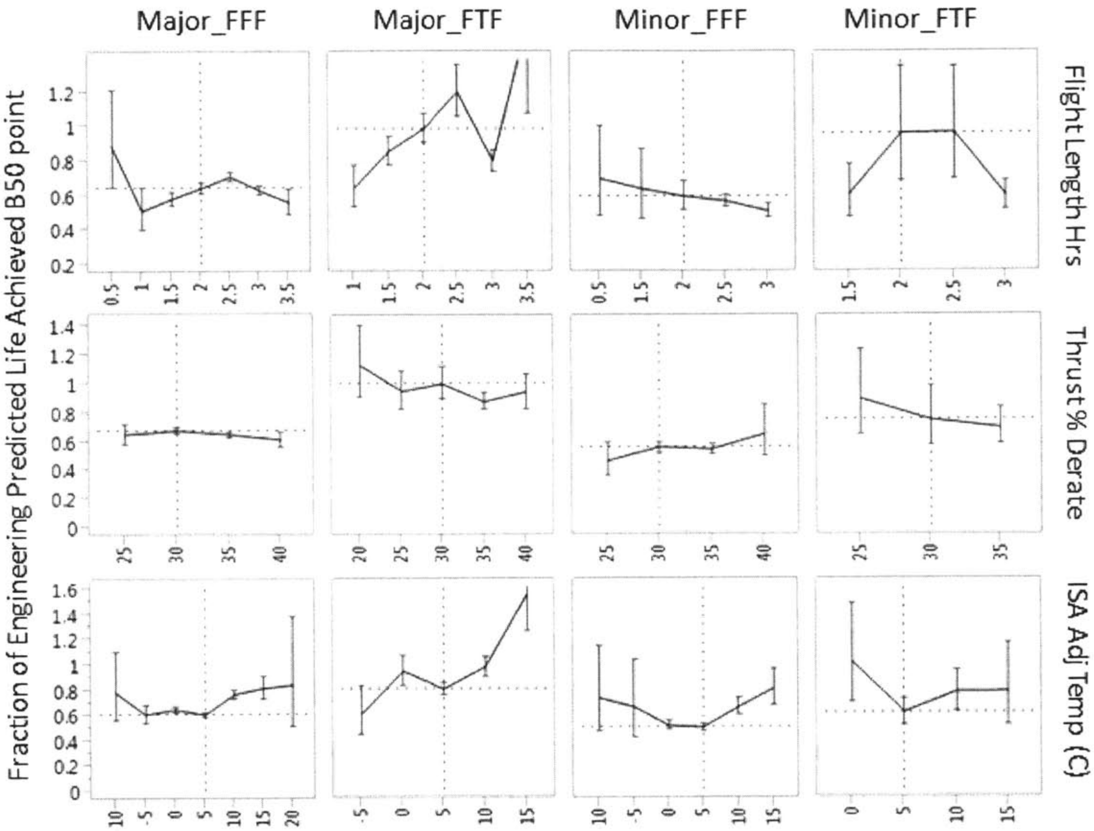

Figure 15: Temperature vs Elevation on interval average level

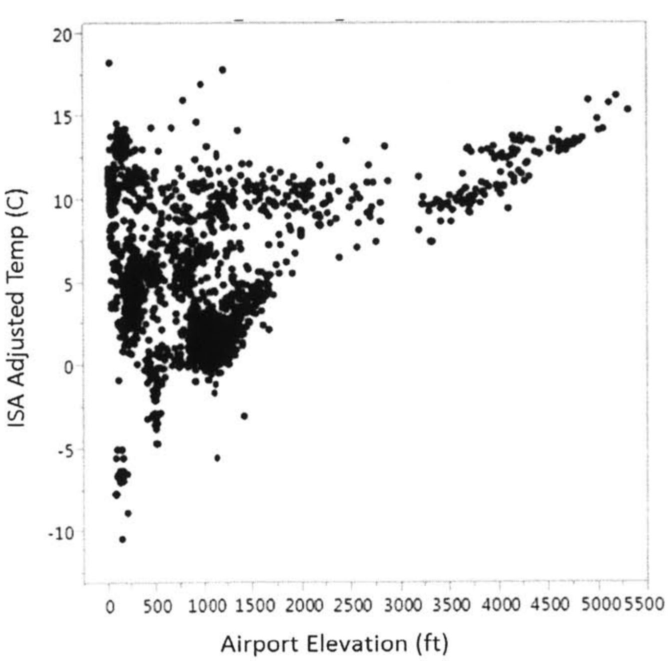

The data does not disprove the engineering model, and it is generated from direct first principles simulation by the Pratt & Whitney engineering team. Therefore the results of the engineering predictions are considered trustworthy as a first order approximation of the effect of ambient temperature, flight length, and takeoff thrust on the engine expected TOW. After correcting for the configuration variance by using the factors in Table 5, we perform a single variable analysis of slope and scale dependency between configuration adjusted life factor (CALF) with results displayed in Figure 16. Analysis is limited to those factors that can be known about the flight prior to scheduling the takeoff. Operating parameters during the flight are not included. The analysis of Figure 16 cannot be done without consideration of variable correlations listed in the appendix. The data is assumed to be fully corrected for temperature, flight length and takeoff thrust. For example, we cannot conclude that 100% thrust reverser usage is good for engine life without considering that thrust reversers are used more actively on short runways, which are correlated to humid airports with low presence of fine grain organic particles. It is possible at this layer that the effect is caused either by the absence of fine grain organic particles or by the humidity difference. Similarly, the cruise altitude effect on engine life is likely correlative rather than causative. Although higher cruise altitude is generally more efficient, short flights require low cruise altitude due to climb time, and long distance flights have lower initial cruise altitude while burning off fuel to achieve higher altitudes. In any case, this work is understood to be correlative in nature and direct action to ensure maintainability of a fleet could not be taken using these data. The data at this stage only provides insight to the potential for higher order effects in each term. Relative humidity demonstrates a possibly quadratic effect which is believable. Humid air reaches a lower exit temperature and requires higher fuel to achieve the same thrust thereby generally increasing engine damage, while at a humidity of 100% rainout will cleanse the air of contaminants. For this reason cross terms including humidity and aerosol density will be included in model construction.

Figure 16: B50 life factor of all engine configurations vs average airport elevation

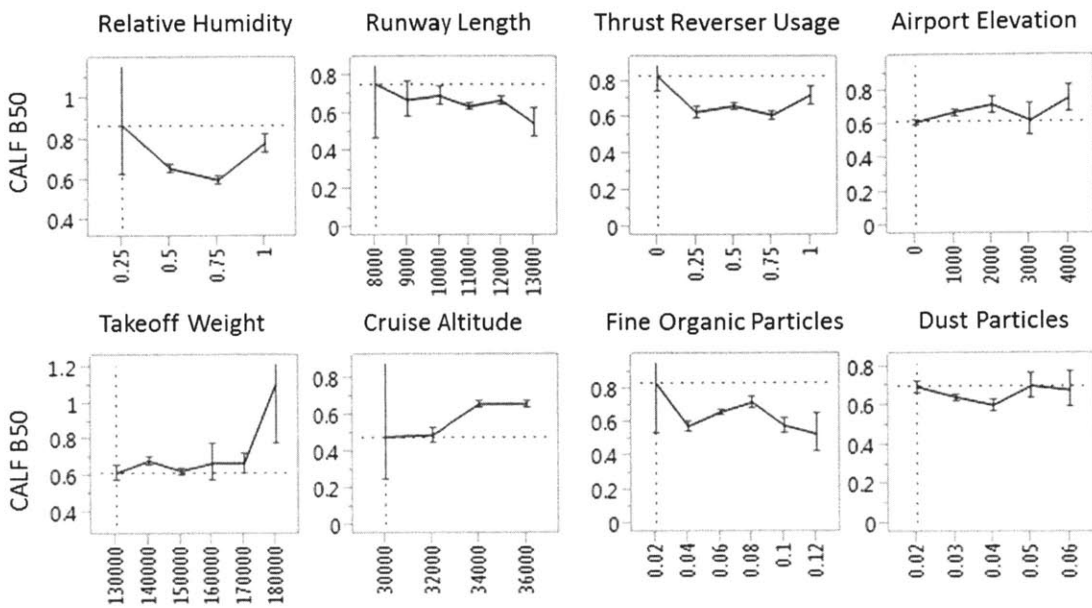

A model is developed to predict time on-wing from parametric survival analysis alone. We consider the effects seen during effect screening to determine which parameters are used in the analysis as linear and which are quadratic. The fit to a loglogistic equation is performed with the following equation:

Equation 3 $$P = \frac{1}{1+e^{\beta*(-\log(CALF)+\varepsilon+\sum \alpha_i * x_i)}}$$

where P is the cumulative failure probability at time given by CALF and $\alpha_i * x_i$ represent all fit parameters and β, ε are fit parameters of scale and offset. The resulting model of best fit contains the parameters and directionality given in Table 6 with Akaike information criterion (AIC) 1320, and Bayesian information criterion (BIC) 1405 compared to a non-parametric fit with AICc 1712 and BIC 1723. The model fits to the training data using a fleet level aggregation. Each airline operator and thrust category represents a single fleet. The predicted distributions of each engine in the fleet are convolved to yield an expected interval for the fleet level. These fleet level averages are compared to historical average interval achieved for that fleet with an $R^2 = 0.22$. A similar model developed using the engine age in cycles at removal instead of CALF and including flight length, temperature, and takeoff thrust disregards the assumptions of the engineering model. This later model yields an improved overall fit with $R^2 = 0.36$. Neither of these models exceeds the existing baseline approach of regional adjustment factors shown in Figure 4. A more complex model for isolation of environmental effects is required that will enable each factor to contribute to the compound model in the way appropriate to that effect. "We just need to be smart enough to ask Mother Nature the right question[47]."

Table 6: Factor contribution to loglogistic model

Factor	Prob>ChiSq	Effect
Cruise Altitude	0.0349*	-
GW Takeoff Weight	0.0019*	-
Airport Elevation	<.0001*	+
Runway Length	<.0001*	-
RH Relative Humidity	<.0001*	+
A1 Fine Organic	0.0004*	-
A8 Total Organic	<.0001*	-
A19 Spherical Dust	0.0057*	-
Thrust Reverser Use	0.0015*	+
Elevation*Elevation	0.1065	-
$A1^2$	0.0001*	+
GW^2	<.0001*	+
RH^2	<.0001*	+

3.5 Multi-model Method

A combination of modeling approaches enables the application of well-informed priors to the model before the application of a general regression study. The problem established requires different model architectures for each layer of information revealed. The proposed Multiple Information Source system identifies three information layers: first, known mission effects, second, censoring and truncation effects, and third, environmental severity. Known mission effects present in the engineering tables convert from the interval duration to the ratio of achieved interval to predictive interval or life factor. Engine configuration laid out in the prior section accounts for major differences in engine life as a result of technology changes or build standard changes at the beginning of the observation period. Censoring and truncation is a condition imposed on the life of the engines by maintenance management policy. Policies for inspection of engines are designated by the OEM with oversight from local aviation authorities. In all operations, the OEM designated policies are followed. However, additional inspections are performed by some airlines due to either company decisions or local aviation authority requirements. When these inspections identify non-conforming conditions the engines are removed for either continuous maintenance work which is targeted to a specific area, or full overhaul. Engines are designed for continuous operation of full interval without targeted maintenance and 70% of engines reach full overhaul without off-wing targeted maintenance actions. The identification of many conditions for targeted work is in many cases subjective, or when objective, the thresholds for tolerance are based upon a planned inspection policy. The acceptable size of grooves in surface coatings on blades is created by engineering assuming that areas will not be inspected again for a known number of flights and is therefore designated conservatively. If the planned inspection policy is accelerated without changing the non-conformance tolerance, then at the next inspection the blade will require removal even though, had it not been inspected, it could have safely remained in operation. For this reason geographic operational area cannot be decoupled from continuous maintenance and the two contribute information at the same level of the model hierarchy. Finally, the condition of the life limited parts and the timing of technology upgrades on the engine and fleet inform the fleet manager in decisions about early removal so that the Major vs. Minor and FTF vs. FFF classifications previously established have significant influence on censoring and truncation.

Fitting engine life factor to a common loglogistic survival plot accounts for the censoring and truncation effects. The decision to include a factor in this level of the model hierarchy is based upon business level knowledge of the daily decision making process regarding engine removal. The second layer of model takes the following mathematical form:

Equation 4 $$\boldsymbol{P} = \frac{1}{1+e^{\beta\left(-log\left(\frac{LF}{1-LF}\right)+a*CM+CF_i+R_j+\varepsilon\right)}}$$

where LF is the life factor, β is the loglogistic shape parameter, $\boldsymbol{CM}$ is the amount of continuous maintenance work performed with coefficient $\boldsymbol{a}$, $\boldsymbol{CF_i}$ and $\boldsymbol{R_j}$ are discrete factors for configuration group and region respectively, ε is the offset of the scale parameter and P is the cumulative distribution probability on which the engine is removed due to deterioration. The standard formulation of the loglogistic scale parameter α can be recovered by the following form:

Equation 5 $$\alpha = e^{a*CM+CF_i+R_j+\varepsilon}$$

The environmental severity is the least understood model with highly covariate effects. Rather than apply this as a parametric effect in the loglogistic model, we isolate it from the censoring by applying it to the p-value of the complete observations.

The distribution range of P requires that future predicted values be on the range of zero to one indicating an advantage to the logistic regression method. P transformed by the logistic equation provides a continuous target for linear regression. Stepwise linear regression is applied using all environmental and mission describing parameters that can be known before the flight. In-flight data may be used to improve model accuracy when monitoring by engine existing fleets, and reduces the flexibility of the model for application to new fleets. Stepwise linear regression will provide a relationship between environmental and mission factors and the expected p-value of the engine when compared to other engines in the total operating fleet. The effect of this third layer of the model is to describe the actual damage severity met by the engine through its operating mission. Reversing this predicted p-value through the parametric loglogistic model applies the behavioral effects of the engine management policy to the underlying damage. This method enables a decoupled treatment of policy effects and severity effects. For example, policy effects of certain airlines in harsh environments are known to result in higher inspection frequency which causes removal earlier than necessary. By categorizing airline policy effects before environmental effects, this bias that policy effects could otherwise incur on geographic environmental effects can be muted.

The hierarchical model structure enables predictions of a fleet expected interval given known mission parameters and environmental values. Several aspects of an airline's operations are well known at the time a contract cost is estimated. The contracts team establishes the airplane mission length, frequency, typical operating temperature, and takeoff thrust within a well-defined range before beginning any contract and during each contract review. While the exact routes the engine will fly cannot be known before contract signature, and frequently change day to day, the list of all airports served by a city with frequency and airframe type enables an estimation of the environmental values that will be seen by the fleet overall. This data is available from the Official Airline Guide(OAG)[51]. An SQL code written in Teradata 14 aggregates a typical environmental experience of a fleet using thrust ratings for airframe configurations, and the city pair service frequency from the OAG. These values are combined with the airlines' known mission plan provide all the information required to perform the estimates in this three layer model.

3.6 Censoring and Truncation

The second layer of the model resolves the censoring and truncation effects. Consider the example of mice in a drug experiment. Censoring refers to individuals in an observation group that have not completed the full observation period and have not yet died. Censoring also refers to those mice that die of causes unrelated to the test. Truncation refers to individuals in the study group that are terminated by the experiment operator after a given length of time. In the engine data example both conditions exist. Engines are either removed for refurbishment due to engine degradation (death), still in operation (censoring), removed for refurbishment due to non-engine causes such as bird strike or scheduling (censoring), or removed due to regulated time limits (truncation). Within the parametric study, censoring and truncation are treated the same mathematically and will be referred to collectively as censored records unless specifically referring to the time limited truncation for unique data applicaitons. The Kaplan-Meier estimator counts these censored observations in the total population (n_i) for all events (i) of deterioration driven removal. The plots used in this analysis are generated using the left continuous formula for the survival estimator $S(LF)$:

Equation 6 $$S(LF) = \prod_{LF_i<LF} \frac{n_i - d_i}{n_i}$$ [52]

In practice d_i is nearly always 1 and LF_i is typically unique to each removal. SAS JMP Pro 11 performs the fit to the parametric loglogistic distribution using maximum likelihood. New construction or full refurbishment of the hot section defines the start of each interval. All observation periods ending in removal are only those removals that incurred a full overhaul of the engine hot section. Early removals for non-engine related reasons include engines removed for service bulletin incorporation, non-engine caused damage such as bird strike or runway debris, and elective early servicing caused by scheduling constraints or customer request.

The third layer of the model aims to identify the reasons that engines fail early as related to the environmental and mission specific factors. When the refurbishment is caused by life limited parts, or when the engine remains on-wing longer than engineering estimation, then the engine has surpassed the planned life and is considered to have completed the observation period. This is comparable to patients of a drug study who surpass the targeting remission time, but are subsequently lost to follow-up. Although the data cannot be used effectively to inform the survival estimator, the fact that it outlived the targeted life provides useful information to the severity model. Furthermore, these engines are strong leverage points. Omitting those engines tips the model results toward only the engines that do deteriorate early without the counter points to the regression model. Model convergence improves on both testing and validation sets when they are included. A significant fraction of the engines used in the regression model are LLP removal instances or engines exceeding 1.0 life factor. Since these items do not have a directly obtainable Kaplan-Meier value, an estimated value is assigned using the loglogistic survival equation of best fit from the second model layer.

3.7 Individual Model Performance

The first model layer applied is a translation from hours or flight cycles into life factor using the ratio of actual length to predicted length using only the engineering provided tables. The accuracy of this first layer can be seen in Figure 4. The previous section detailed the construction of the second and third model layers. This section reviews the performance of these models, insights from plotting effects globally and how the whole system is applied within Pratt & Whitney.

3.7.1 Second Layer - Parametric Survival

The second model layer performs a parametric loglogistic model with a target of life factor, and censoring of non-deterioration related refurbishments. Study of the industry-wide behavioral effects that influence the censoring and truncation yields three factors of primary influence which are used in the model: total off-wing targeted repairs referred to as continuous maintenance (CM), region of operation split by governing authority types and continental area (R), and configuration including both technology level in the engine and build plan of the life limited parts (CF). Categorical parameters enjoy one degree of freedom per category. Equation 4 is fit using SAS JMP Pro 11 using equal weighted error and chi-square tests for coefficient values. Of the engine intervals used in the model approximately 40% were uncensored values and 60% were right censored values. The model is built three ways to validate the assertion that the three factors selected contain the majority of information encoding possible at this layer. A non-parametric model provides the worst case AIC and BIC while a parametric model using all known environmental and operational factors provides the AIC and BIC for an overfit condition. The latter model is similar to the model performed presented in Table 6 except that all factors are included rather than a limited set and the Life Factor is used as the target rather than the CALF. The performance of AIC and BIC for the three models shown in Table 7 demonstrates that while adding factors beyond the first three improves the direct fit of the model reflected in a decreasing AIC, BIC which is adjusted for degrees of freedom does not move significantly.

Table 7: Parametric Survival Model compared to extreme case under fit and over fit methods

	Zero Factors	Three Factors	All Factors
AIC	875	332	188
BIC	886	401	355
DF	0	10	27

The three factor model passes a chi squared whole model fit test with probability <0.0001 and all three factors individually pass with probability <0.0001. The coefficient directionality matches those expected from Table 5 and Figure 12 and the regional effects are shown in Figure 17. The resulting model is applied to all engines to convert from life factor into p-value which represents the survival positioning of the engine relative to all others after adjusting for the three factors used in the model.

Figure 17: Regional effects on interval life factor B50

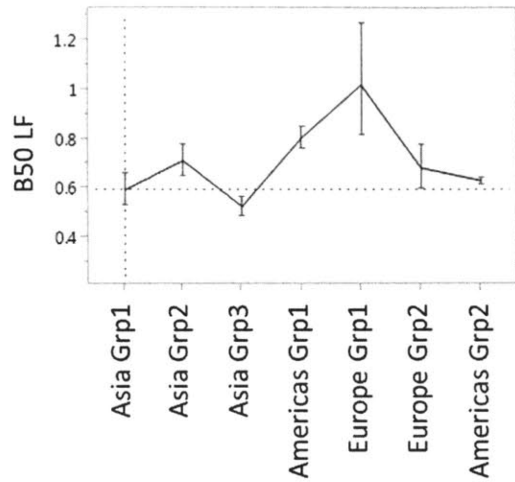

3.7.2 Third Layer - Logistic Linear Regression

The third model layer performs logistic regression on a target probability from the second mode layer with regression factors which characterize the environment and mission conditions of the engine throughout the interval. Average values and binned ranges used for each environmental variable allow for both linear and non-linear effects without the use of higher order terms. The binned ranges are split to three regions in each parameter of high, medium and low where the value recorded is the percentage of flights in each category. One third of the data is withheld by random assignment for validation of the model.

The logistic linear regression treats each engine interval as a single point in model training. Engines with higher numbers of flights represent more information about the effect of the environmental factors than engines with very few. The regression model error calculation is therefore weighted by cycle number so that high leverage data points with thousands of flights are given proportionally more sway over the parameter estimates than the engines that were removed quickly. There exists the possibility of misclassification in the data. The parametric survival estimation methods are not sensitive to a small number of misclassifications. The early removals represent strong leverage points in a logistic regression model and the use of cycle weighted error calculation minimizes the risk of this effect. The impact of this decision will be explored later in the chapter.

The following three parameter pairs exhibit high correlation of estimates and the former term in each pair is suppressed from the stepwise regression model: Average Gross Weight to Average Takeoff Thrust

Derate; RH*A14 High to A8 High; Average A19 to Average A2. The final model includes parameters in multiple ways. Aerosol effects with cross terms of humidity pose a challenge to simple interpretation of the results. Table 8 shows the standard beta which is a normalized version of the factor coefficient and enables more effective comparison of terms to understand the underlying effects. Here negative directionality results in lower p-values and therefore lower life factors. The directionality of Thrust Reverser usage rate aligns with the expectation that using high thrust during landing should have some detrimental effect on engine lifetime. The idea that cleaning actions are negatively correlated with life can be understood when considering that most airlines perform the cleaning only in harsh environments resulting in a direct sampling bias error which is caught by the introduction of cleaning action as a factor. Unfortunately this incurs a bias that makes it difficult to understand the true effect of the environment. If cleaning action policies remain largely unchanged then the model will continue to represent reality. However, the model cannot be used to allocate a shadow price to such cleaning. The appearance of Flight Length and DTamb in the final model indicates that the first principles simulation by engineering either does not fully account for the effect of these parameters or that there are other effects not being sufficiently included by the factors selected. The effects of aerosols and humidity will be analyzed in the next section.

Table 8: Logistic regression factors of primary model

Factor	t Ratio	Prob>|t|	Std Beta	VIF
Intercept	0.36	0.7219	0	.
TR Usage Rate	-1.26	0.2081	-5.8	2.67
Cleaning Action	-3.58	0.0004*	-13.3	1.70
Flight Length	4.12	<.0001*	18.1	2.38
Cruise Altitude	-2.24	0.0255*	-11.9	3.48
DTamb	3.25	0.0012*	16.3	3.13
A14 Mid	3.60	0.0003*	14.2	1.93
A21 Mid	-4.58	<.0001*	-22.4	2.98
RHxA3 Mid	-2.33	0.0203*	-14.0	4.51
A1 Low	-3.25	0.0012*	-19.5	4.47
A8 Low	1.26	0.2071	6.1	2.90
Elevation Low	-4.03	<.0001*	-21.8	3.63
Elevation High	-7.36	<.0001*	-43.3	4.28
RHxA1 Low	2.38	0.0174*	11.9	3.10

Factor	t Ratio	Prob>|t|	Std Beta	VIF
RHxA2 High	-2.57	0.0103*	-20.4	7.80
RHxA8 Low	-2.81	0.0051*	-19.0	5.67
RHxA14 Low	1.96	0.0498*	11.9	4.56
RHxA21 Low - Dust	-2.20	0.0284*	-13.7	4.81
A19 – Dust	1.95	0.0517	13.7	6.10

The logistic model exhibits a relatively low adjusted R^2 of 0.312, but does pass lack of fit tests with an F Ratio of 43.4 or probability 0.0047. With the high number of variables and expected degree of variance within the data neither of these figures is necessarily compelling. A comparison of the adjusted R^2 to the R^2 of the validation data of 0.266 reveals that the model is at least detecting features that exist across the population.

When the final model is completed and all three layers combined, the results at a fleet level compare favorably with the baseline model presented in Figure 4. The final R^2 of fleet level predictions is 0.7966, and 80% of fleets fall within the predicted range for average interval when at least three engines in the fleet have uncensored removals. This compares to the regional method R^2 of 0.526 with 66% of fleets falling within the predicted range. The mean square error of the regional method is 13 million in flight hours, while the complete three layer model results in an MSE of 6 million. Furthermore, when the regression is repeated including module level performance data such as internal temperatures, fuel flow and vibration levels, the adjusted R^2 declines and MSE rises. More pertinent to the application of the model, though, are the insights provided by the derived effects discussed in greater detail in the following chapter. We argue that these observations may be applied to further fleets outside the test and training data used in this analysis. Despite the usefulness of the model, it is important to limit assumptions about the nature of these relationships as we recall the assertion by Box and Draper, "Essentially, all models are wrong, but some are useful[53]."

Figure 18:Total model fit to fleet level Weibull derived mean intervals

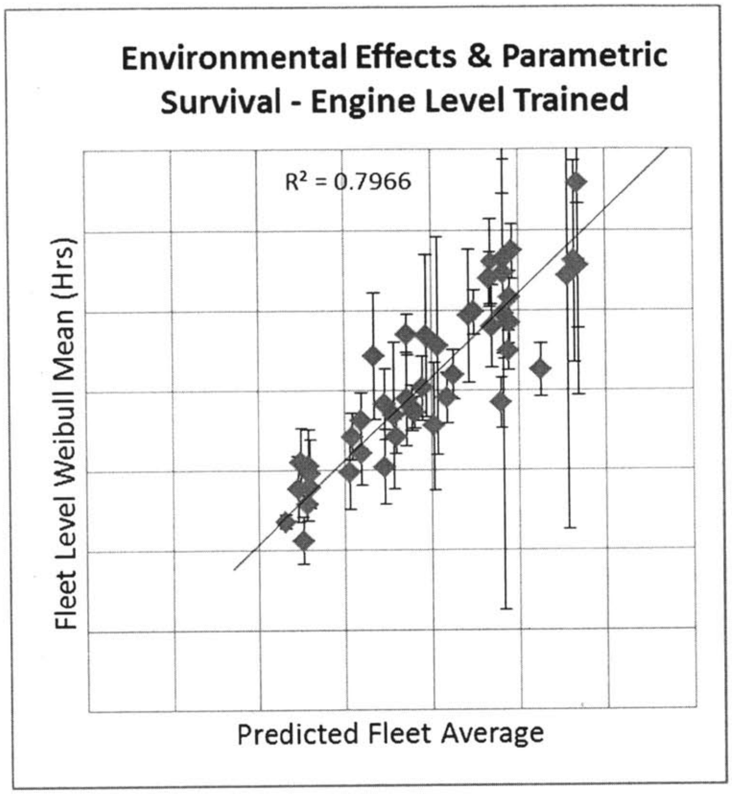

3.7.2 City Categorization

Data visualization on a map overlay enables quick validation of the system. We aggregate all airport data to a single row by averaging environmental values over a two year period. Assuming a standard length mission at 35,000 feet with a typical aircraft for the studied fleet provides the necessary information to fill all necessary factors of the predictive model for a single city. The resulting model provides a single severity score for each city. The formula for adjusting air temperature to ISA involves elevation, which is one of the parameters in the severity model. Therefore when visualizing the effects of the city and region based severity, the temperature aspect of the first layer model needs to be applied as well. In order to estimate temperatures properly, the list of airports used includes only those airports with flights of engines where a temperature reading from the engine could be used instead of airport weather data. This adjusts for the fact that some operators avoid flights in the hottest parts of the day. The final severity for each city plotted onto a world map from NASA Terra provides a method of confirming that regional effects are well behaved. Figure 19 displays the life factor for each city using the method described with 2012 and 2013 environmental data. This data presentation is largely isolated from the effects of flight length and takeoff thrust and represents the effects on severity caused by atmospheric and airport specific factors only. Two primary effects are displayed by the map. First, the effects of temperature tend to dominate throughout the map in that equatorial regions tend to be more severe than polar regions. However, the effect of local atmospheric factors is sufficiently strong that any two adjacent airports are likely to be significantly different in severity. This is confirmation that the sub-continental regions used in the second model layer do not dominate the system. There is no apparent border visible between regions.

Figure 19: Global Severity Map

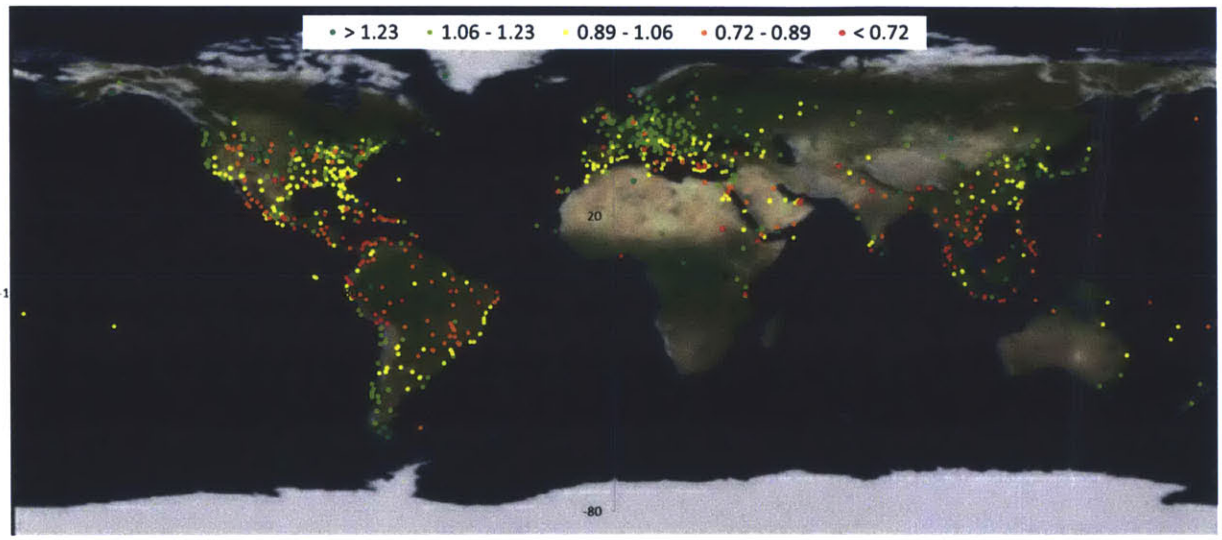

3.7.3 Total model sensitivity to methods

The use of cycle weighting for the environmental logistic regression model enables the model to successfully pick up on meaningful information in the data. As described previously, the model has an adjusted R^2 of 0.312 on training data, and an R^2 of 0.266 on the testing data when one-third of engines are withheld. When the same model factors are selected and coefficients are scored with equal weighting per engine, the adjusted R^2 on training data and R^2 of testing data drop to 0.135 and 0.037 respectively. If the forward stepwise regression is executed again and collinear terms are pared down as done for the primary model, the adjusted R^2 on training data rises to 0.283 and the R^2 on validation data is 0.122. When the model is trained on equal weight engine intervals it is being dominated by leverage points which had low exposure time to the actual environmental factors. Flight cycle weighting represents an important insight critical for creating a functional model that depends upon a clear understanding of the underlying structure being described by the model.

3.8 Application to Current Business Methods

The system of models developed here provides three direct benefits. First, as the system's principal goal, it provides an estimate for fleet level average time on-wing. Second, the system provides a shadow price for a number of decisions including maintenance actions and fleet management decision. Third, engine level predictions are developed for the active fleet enabling better planning decisions.

The system is easily operated and relies upon a well-integrated data layer. SQL code contains the business intelligence of combining multiple source systems and aggregating engine data into a standard input format for analysis. Daily or weekly execution of these codes is capable of maintaining the infrastructure. SQL code applies the conversion of the first model layer of interval length into life factor by referencing interpolated engineering tables maintained by ES. A data stream in SPSS Modeler 16 retrains the second and third model layers quarterly, or when new environmental data is acquired for model improvement. OAG data on Teradata with planned flight frequency by airline and airport combined with environmental data provides a table of fleets with best expected environmental values in the same format as the training data. The SPSS data stream pulls the reformatted OAG data, runs it through the second and third layers of the model and provides a table of fleet level predicted life factors back to the Teradata server. This table may be updated frequently, but in practice updates are only necessary when the fleet operations for an airline change significantly, or to reflect new environmental data downloaded from NASA Terra.

An Excel based interface consumes only the model output for standard use. In this way, fleet predictions may be simplified to the same level that they were previously when applied by the business with minimal retraining required and enabling full flexibility of staff using the models. The predicted time on-wing (TOW) obeys the following equation laid out in this hierarchical model:

Equation 7 $$\boldsymbol{TOW = Hrs(FL, PDR, CA) * TF * LF * CF}$$

Where FL, PDR, CA are flight length in hours, percent derate, and cruise altitude respectively, and TF, LF, CF are temperature factor from the associated engineering table, predicted life factor from the fleet severity models, and configuration factor. Configuration factor is added here to enable the use of the model for fleets that do not have sufficient data for the configuration to be included in the model directly. When technology upgrades are made, an informed declaration about the time on-wing effect years before sufficient operational data may be used to validate the claim using lab testing and similar engines as was demonstrated as early as 1980 by Pratt & Whitney's expansion of use of single crystal blades[54]. Adding this feature to the model enables the business to perform these flexible adjustments to the prediction based upon non-modeled knowledge.

4.0 Validation and Implementation

4.1 Benchmarks

The proposed system is challenged by review alongside two important reference points. First the airport level effects are compared to published effects observed by a competitor in the aerospace industry. Although the models are developed from a different basis, the general trends between them provide insight as to the stability and value of the proposed model. Second, the model is compared to global trends in temperature and aerosol levels. The potential causality of these statistical effects is considered along with a discussion of related effects that may be leading to confusion of terms. The observed effects are seen in directionality and geographic location to previously published work, although new effects in Caribbean airports are subject to some uncertainty.

4.1.1 External Benchmarks to Competitor Publications

Work presented by General Electric (GE) in 2011 on the topic of regional severity provides a meaningful comparison for the current results. A side by side comparison is presented in Figure 20[13]. The current model exhibits a high degree of variability city to city which is described in greater detail in the following section. In addition, the current model shows the severity calculation only for those cities where engine takeoffs have provided sufficient real flight temperature data for a schedule adjusted severity. On contrast, the smooth transitions of the GE model appear constructed from global average temperatures and are similar to Figure 21. On the whole, the GE analysis tends to cluster the same cities into similar categories as does the current model. This is likely due to the overriding effect of temperature in the tropic cities. The dust impact categorization varies between the two models. It is not clear what source was used for the dust study of the GE model, nor can any numeric severity be acquired for quantitative analysis between the models. The regions of greatest discrepancy appear to be Middle East and possibly Europe, although the GE chart is obscured in that area. Within these regions to fleets have sufficient data in the testing set to conduct a Weibull based average interval analysis for comparison to the current model. Although both of these fleets are within the 95% confidence band of the current model's prediction, they observed average intervals of 95.5% and 88.5% of the models' expected average. This is a possible indication of a geographic error in this region by an order of 5-10% which could be caused by insufficient interval data in the region.

Figure 20: Comparison of city environmental categorization by General Electric to the current model

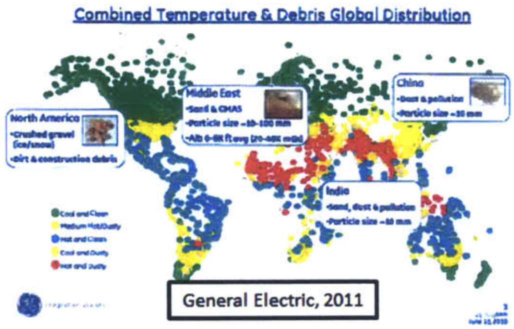

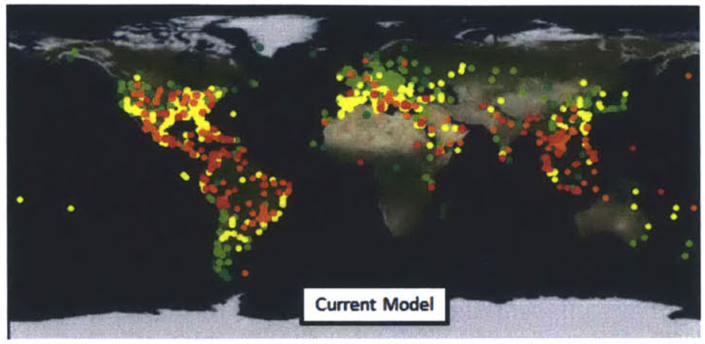

4.1.2 Comparison to overall Environmental Behavior

The results of the model may be compared visually to the global distributions of its key geographic drivers: Temperature and Particulate Matter concentrations. The observed temperature severity compared to the global temperature distributions shown in Figure 21 shows a strong overlap with two notable exceptions. First, high elevation airports along the Andes and Himalayas are both penalized by the model beyond the degree that the ambient temperature would predict. Secondly, mid-continental hub airports, notably in the United States, display both high severity and high variance. Within the data this is observed to be an artifact of the flight schedules. Airlines servicing mid continental cities in the United States make use of mid-day connecting banks, with traffic peaking in early afternoon[55],[56]. This results in a flight sampled temperature significantly above the average experienced by the airport. For example, although typical airports within 500 km of Denver (KDEN) have an annual average ISA adjusted temperature of 2 degrees C, the ISA adjusted temperature at Denver is 8 degrees C. This effect is also seen in adjacent airports serving different traffic types. Los Angeles has an average annual ISA adjusted temperature of 3 degrees C. The international airport (KLAX) which serves longer stage length flights morning and evening experiences an average takeoff temperature of 2 degrees C while the temperature at neighboring KLGB servicing shorter stage length traffic during business hours is 4 degrees C. This causes a net 4% decrease in time on-wing at KLGB vs KLAX based solely upon flight schedule.

Figure 21: Global Temperature Distribution with Airport Temperature/Elevation Severity Overlay[57]

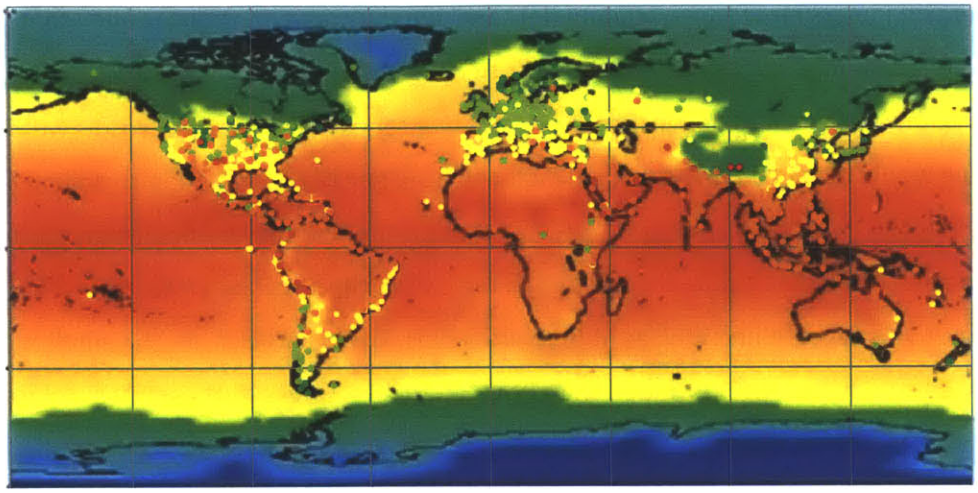

Just as the elevation and temperature parameters cannot be considered independently due to their natural interdependence, the effects of aerosol and humidity must be jointly analyzed. Where high humidity is observed, aerosol levels will naturally be suppressed by rain-out and atmospheric models tend to show high correlation between suspended aerosols and cloud formation[58]. The model parameters shown in Table 8 show a clear correlation of terms. When the cross terms between relative humidity and aerosol levels are removed, the model cannot converge above an adjusted R^2 of 0.15. Figure 22 shows maps of water column[59], and PM2.5 concentration[60] alongside the aerosol effects and demonstrates a negative correlation between humidity and PM2.5.

Figure 22: Comparison of Global water column and PM2.5 concentrations to combined severity effects

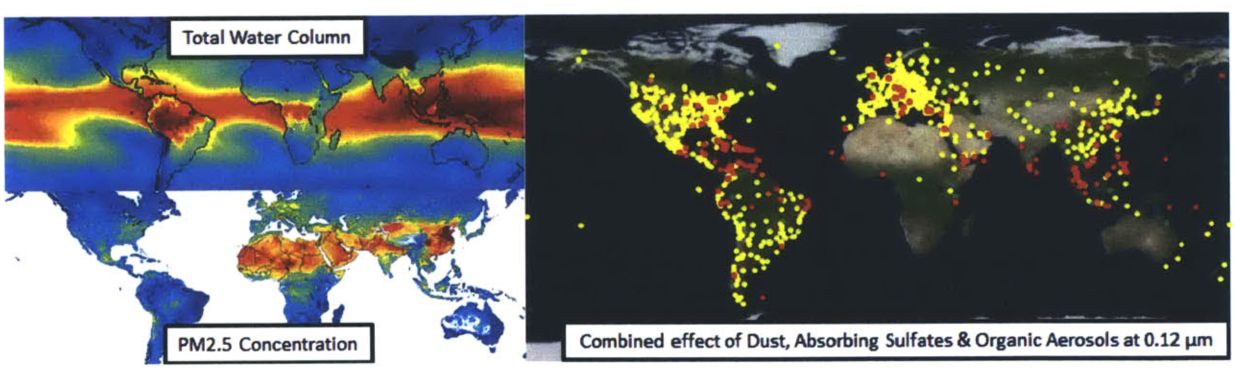

Lugo-Leyete et. al. have shown that the turbine inlet temperature will decrease with rising relative humidity for the same air volume due to increased thermal mass of the air[61]. Aerosols are suspected to have general corrosive and abrasive effects which contribute to engine damage, and therefore a negative correlation between humidity and aerosol complicates the intuitive understanding of the model. The aerosol factors picked up by the model conclude that large particle size pollutants, such as coal dust, which is characteristically 2.0 μm, are less significant than smaller absorbing particles. These particles include organics such as pollens, fine soot and fine sands. If ocean salt aerosols were detrimental, we would expect a non-absorbing factor to have negative correlation to engine life, but this is not observed. Of interest geographically is the fact that northern China, although high on the PM2.5 map, is low on aerosol severity, while the Caribbean region has the opposite effect. This is led by the model training on A14, absorbing particles at 0.12 μm radius. These particles are too small to be included in PM2.5 and have longer duration than PM2.5 particles. Larger dust particles included in the PM2.5 value will not tend to reach the Caribbean from Africa, although MISR data does indicate elevated values of total spheroidal dust (A21) in this region. Figure 23 displays these two particle concentrations showing that although particle size of A21 may have decreased, the overall quantities reaching Caribbean airports remains fairly high. At the same time, the dust reaching South American airports from Africa is elevated only along the coast, while inland regions are benefited by rain-out. While the data source for the current study and the PM2.5 readings is the same satellite source, the current study has focused only on readings within 30km of an airport with engine service resulting in significant local data dispersion. This dispersion is expected to cause the elevation of A14, too small to appear on PM2.5 ratings, and generally sourced by local environmental effects such as burning biomass, pollens and certain types of city pollution.

Figure 23: MISR average readings for 2012/2013 A14 and A21 within 30km of selected airports

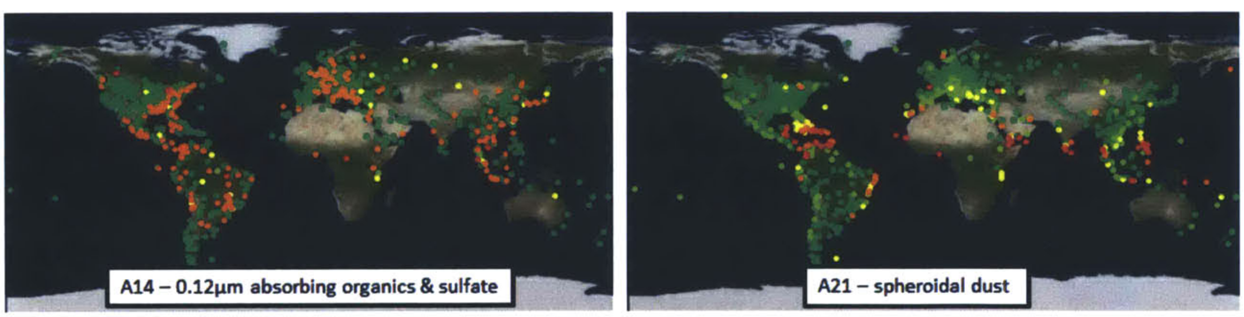

4.2 Organizational Analysis

Even a perfect model must be applied to a business solution by an organization which incurs its own effects and constraints. An organization's structure, politics, and culture each have a significant effect on the performance of the business. Alignment between these three aspects and the business goals is a required underpinning of any project's successful implementation. The following section describes the current state of the Pratt & Whitney aftermarket and IT organizations and the interaction of sub-teams in the company referred to as 'business units'. The data acquisition process and model development process are reviewed in detail as to the interplay between organizational structure and organizational performance. Finally, recommendations are made regarding the optimal data management structure between both the IT core team, the individual business units, and the proposed Data Analytics organization.

4.2.1 Overview of Organizational Effects

Pratt & Whitney aftermarket is aligned as an independent Profit & Loss center (P&L) with a number of core functions such as IT aggregated to the Pratt & Whitney corporate level. The incentives and P&L splitting within Pratt & Whitney have changed over the recent decades with a concerted effort of P&W to aggregate P&Ls to leverage synergy[62]. The recent acquisition of International Aero Engines (IAE), which contributed the majority of the engine install-base to the aftermarket support, enabled the company to consolidate its relationships with narrow-body customers[63].

The structural alignment of the aftermarket division is the result of the collection of many independent P&L's[64]. Pratt and Whitney has a made concerted effort on consolidation of overlapping business services following this collection and the IAE acquisition[65]. IAE and P&W independently sold and supported engines for over ten years prior to the integration. All fleet support mechanisms for structure, systems, and customer relationships were independently formulated and integration of them to realize synergy without disrupting the customer support requires careful attention and time. Data analytics at Pratt & Whitney represents a central aspect of this integration effort[66]. The structural differences in organization extend into the data for this team as well. After the IAE acquisition, some of the engine's on-board data systems used to support the fleet are split between the Rolls Royce system OSYS[67], and the Pratt & Whitney system ADEM[68].

Pratt & Whitney cultural effects in the engineering department are well documented by Bowen and Purrington in 2006, and are observed to be consistent in 2014[69]. The airline industry is infused by an overt priority on safety for the flying public[70]. This results in a prevalence of cautious planning, and a sense that things cannot be rushed. Third party auditors are a familiar part of daily work and awareness of regulations and authority pervade the working culture.

Bowen and Purrington observed that political power within Pratt & Whitney is held both along the traditional hierarchical lines, and by senior engineers. Within the aerospace industry, a majority of employees have more than 10 years' experience, and it is generally observed that experienced staff requires lower direct involvement of the management team. The resulting high number of managers in the aftermarket division leads to a dilution of political power from the 1st level. Financial services and insurance companies tend to have lower management ratios in the range of 8:1, while technical support centers have been observed to be in the range of 5:1[71]. These figures would be fitting targets for aerospace aftermarket organization whose functions are a mix of financial, insurance estimation, and technical support. With lower management ratio, and higher average experience level, political power will tend to be shared between senior engineers and the hierarchical management system. Senior engineers with significant tenure in position, referred to as Discipline Chiefs by Bowen and Purrington, are regarded as the owner of a given set of operational practices. This position is typical of the aerospace industry and is symptomatic of an established labor force with low turnover. Indeed, the AIAA estimates that 50% of US aerospace engineers are currently eligible to retire[72]. This makes documentation and interchangeable labor a critical part of business sustainability in the near future across the industry. Bowen and Purrington observed that the reliance upon discipline chiefs contributed to the difficulty of implementing documented work practices. Decreased documentation tends to increase the cost of rotations and training and decrease organizational repeatability. At the same time, the role of discipline chief was observed to enable the promulgation of best practices to the engineers in the discipline chief's field. In the early 2000's, Pratt & Whitney underwent a number of documentation efforts to enable effective knowledge transfer by written word and ownership of updates to these procedures was considered the most tenuous risk[69]. Discipline chiefs, many of whom were also technical fellows, exercise greater influence on the organization than first level managers. The result of this high degree of knowledge based power was a general aversion to rotations within the company. Employee rotation was feared by many managers due to domain critical knowledge and a lack of proper documentation for supporting transfers. To counter these effects Bowen and Purrington had recommended new rotation initiatives. However, in the opinion of this research, this work did not go far enough.

The importance of employee rotation programs is perhaps best displayed by the SOX-mandated 5 year limit on auditors. A number of methods for mitigating knowledge loss are presented by Sanders, Steward, and Bridges[73]. These methods are aimed at making the rotational planning a continuous effort with improved and sustained documentation and could be adapted by large aerospace companies to facilitate career mobility, improved knowledge documentation and new idea development.

4.2.2 Data Acquisition & Data Maintenance

Pratt & Whitney data maintenance is a shared effort among customer facing business departments and internally facing information technology (IT). Each business unit contributes and references data continuously, and a naturally developed dispersion exists in the level of integration between an individual business unit and IT. Within the archive, data and preserve functions, Pratt & Whitney IT relies regularly upon 3rd party IT partners through the use of stable long term contracts[74],[75]. Ingest, Administrative and some Data functions are retained by Pratt & Whitney. Access functions are generally left to the business units, with oversight and support from the Administrative and Data functions as required.

The Pratt & Whitney data systems employed for this work compare generally well with World Data Systems on as measured by the WDS survey proposed by Laughton[44], although for reasons of confidentiality the survey results of the P&W data systems are not presented here. Table 9 shows that the functional scores are at or above average in four of the six functions. Two functions stand out across the industry as areas of opportunity for improvement – Preservation and Access. The key reason observed for higher scores on Ingest in successful companies is a result of a pre-ingest function similar to that proposed by Laughton and du Plessis[43]. These pre-ingest domains had generally developed from necessity at the business unit side. Certain critical data is pre-filtered, audited and formatted for upload by engineering staff prior to submission to the IT data systems.

Table 9: Average World Data System Scores by Segment

	Ingest	Archive	Data	Admin	Preserve	Access
WDS Avg	76%	66%	79%	65%	63%	54%

The consolidation of business and data function mentioned previously with regards to the IAE acquisition may be expected to lead to improvement in the Preservation score. As the integration of older data systems completes, any company will improve its ability to preserve data, although care must be taken to monitor technology and formatting to ensure availability of future staffing. A majority of data architecture decisions for the preservation and integration of data structures should generally be led by the business unit engineering teams with data relational knowledge. Education of these engineering teams regarding proper data curation methods is critical to the improvement of long range data management. This may be done through close integration with local colleges and universities. In the fall of 2014, Pratt & Whitney launched an initiative to provide accredited coursework for a master's degree in data analytics through a local college to all engineering staff. This type of action will reduce the risk of data obsolescence and increase the probability of data use within the company. Some members of the organization had already enrolled in similar programs due to job requirements for technical master's degrees. Others who had met those requirements through different degrees felt that the work to attain a degree was unnecessary and preferred taking targeted coursework or seminars[76]. The deployment of continued education may take the form of formal coursework, or developing a relationship with academic professionals who study with the company engineers. The former will supply expert knowledge to certain team members, while the latter will provide a source of ongoing research. This is distinct from hiring a Ph.D. level work force that integrates with the existing staff. The unique vantage provided by active professors is their outsider perspective.

Studies have shown that high grade level engineers are the least likely to attend courses, even when courses are freely available and encouraged by their employer[77]. Continuing education is of the most importance for these engineering leaders due to their influence in technology decisions. Pratt & Whitney leadership is encouraged to attempt both methods of continuing education concurrently. It can be expected that continued education regarding data management among the IT and Business unit engineering teams combined with initiatives currently underway to centralize data storage will improve the Preservation score of Pratt & Whitney.

One senior engineer working in aftermarket commented toward the end of this research project "Where did you find all of this? It's amazing[78]." The data used in the project was generally available to any Pratt & Whitney employee; however, access to it and understanding of it was granted individually by business unit domain owners. The IT data function consists largely of maintenance of data schemas and maps of relational tables. For an understanding of which data types were equivalent between databases, interviews with business units experts are required. In the absence of company-wide advertising of available data systems, finding these experts requires personal requests. As a result, when data is required across systems between two groups, it is tempting to replicate the data in full on each side. This enables each system owner to maintain full knowledge of their product, without dependency upon other groups for system maintenance. The resulting downside of this solution is a risk that some teams will not know the origin and context of the data and may act on it without full understanding.

The delegation of the Access function to the business units is not entirely out of place. Maintenance of complex queries that supply a business intelligence layer over the underlying data architecture relies upon both firm knowledge of the server language and knowledge of the business process being served.

Furthermore, advertisement of the available data and cross-functional integration both require integration with the business units to identify data needs. The advantages of centralization which gave rise to the development of a central IT role in most organizations remain relevant and unmet in the current state. The development of a centralized business unit responsible for data curation oversight from a business perspective may serve to fill this gap when it is done in concert with the existing IT central structure.

4.2.3 Model Development & Deployment

For a model to be successful it must present value, flexibility and relatability to the organization. Proper development of predictive models requires both access to and knowledge of all pertinent data. Thus a data analytics organization would ideally contain members from all parts of the organization. This provides the group with not only a clear insight to critical business needs, but meaningful experience in determining the value of a project to the company. The ability of a model to be relatable is a ready concept to the engineers on the maintenance cost group who have spent years working with existing data and converting the pure analytical results into a format that can be documented into a sales contract. The constant changes in these contracts and customer requests drives the flexibility of the organization to ensure responsiveness.

The greatest benefit to the data analytics team, its functional excellence, is also its largest risk. As the team develops in experience, the tendency to assign a discipline chief will rise, and the organization may become its own new silo. This might enable new modeling methods to be deployed, but is more likely to result in an aversion to job rotation as previously discussed. This will eventually lead to a stagnation of the extensive cross-functional experience within the company that the team currently has. For this reason, relying upon external contractors to provide data modeling insight may be a prudent decision for many groups. IBM and other data analytics partners provide P&W with on demand staff highly skilled in analytical methods[66]. While data analytics talent should be developed internally to a limited degree, there exists a risk that professional opportunities for rotation and advancement could be limited within any engineering or manufacturing company. This could lead to stagnation, or turn-over. It is critical that the engagement between the data analytics team and any external partners occur at the lowest possible level and work through individual data decisions as outlined in the system dynamics section of this research. The partnership had originally aimed for the following simple boundaries: company provides data, partner provides model, company uses model.

Figure 24 displays a more balanced approach that better captures the internal knowledge of the company regarding the underlying structure of the problem being modeled. The proposed framework is understood as a continuous improvement cycle with continuous review at a detail level. One complaint regarding prior work with analytics partners was present throughout the company, "How are we going to do this? These meetings are high level[79]." Under the proposed framework, when a new improvement in data presentation or modeling is conceived by the business user with the data analytics team, an exploratory analysis is first performed in conjunction with engineering experts on the topic. The key output of this phase is a system dynamics view of the problem being described which should contain both social and technical factors influencing the model. As is shown in the current research, systems level analysis is important to enable the proper identification of data sources and prioritization of that work. With a broad system definition of scope to the project, funding can be supplied to gather data sources and explore applicable modeling techniques. At the early systems model phase, both the IT partner and Data methods partners are engaged by the analytics team which is also responsible for the development and deployment of the draft model. Once any new model is deployed a pilot phase is required prior to full implementation to validate the assumptions and ensure that existing business processes were not overlooked. The continuous monitoring of the model focuses on areas where the model or system is incorrect. These errors do not cause suspension of the model, but rather inspire a new pass through the development cycle to continuously improve the model.

Figure 24: Proposed Model Development Framework with Overlay of the Deming cycle

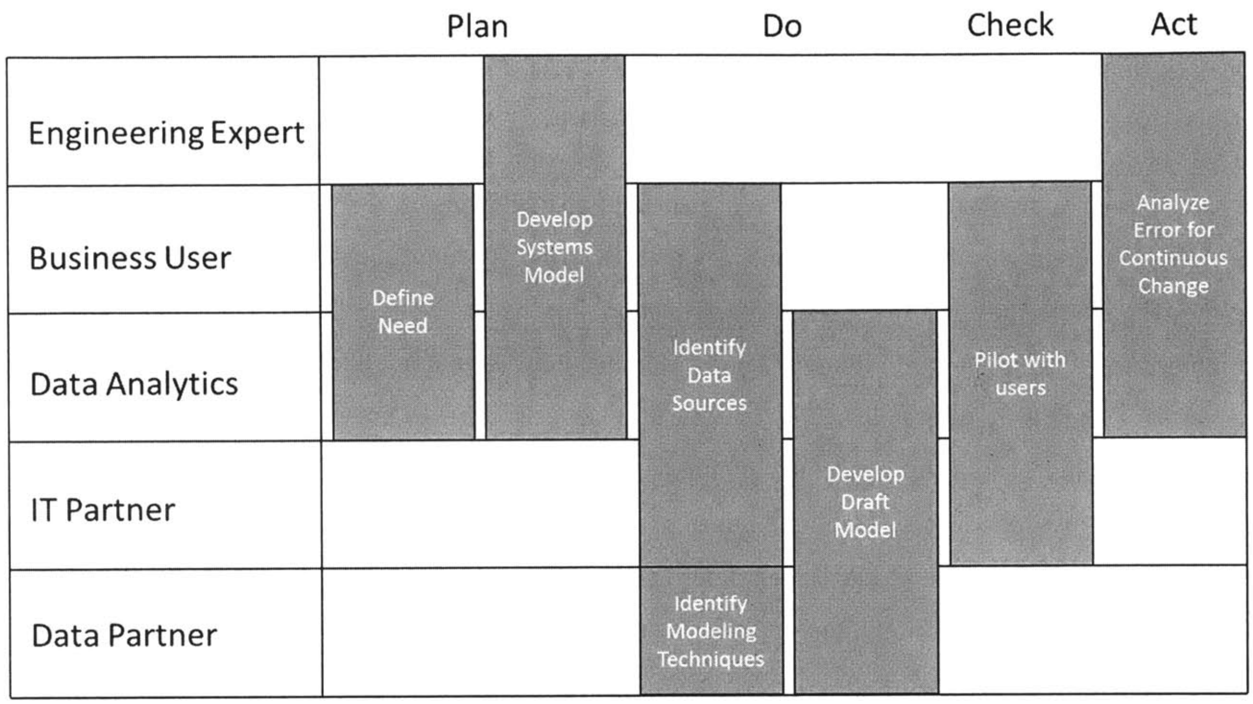

4.2.4 Recommendations for Governance and Structure

Quantitative research on organizational structures has shown that dynamic businesses with a high degree of variability between customers benefit from structural alignment by customer. At the same time, further improvement has been seen in organizations able to provide asymmetric support where larger customers received higher staffing flexibility[80]. This is possible in organizations that categorize staff with deference to process dependencies allowing pooled resources for asymmetric support. Based upon these findings, the first hierarchical split of the data analytics team is designated functionally to maintain functional disciplines and also align with existing corporate culture. The second organizational alignment proposed is a customer alignment through each function of the organization. This lower level split in the teams enables specialization within the function addressing the types of questions commonly received. Figure 25 demonstrates the proposed organizational structure.

The Data Curation team is included within the business unit with functional overlap to the existing IT structure in the three areas of Ingest, Data Management and Access. From this positioning in the business units directly supporting model development, the Data curation team will be capable of developing and successfully advertising the company data from a vantage point generally not available to IT organizations. This team requires close integration with IT, and is intended to be collocated with IT staff in the parallel functions, although the functional alignment remains within the Data analytics team. It is important that this group rotate frequently through positions in the business side of the company to both remain fully effective as advocates of the business. Rotation will also ensure that they continue developing their engineering skills throughout the aftermarket support division, decreasing the risk of becoming an extension of IT in both function and skill base.

The Model Development team performs major oversight and guidance of new model development. The expertise required by this team is not specifically mathematical rigor, but system dynamics. Familiarity with modeling techniques is a key aspect of the function, but the strength to the company is in combining those techniques with the engineering domain knowledge unique to the OEM. The Model Development team is therefore aligned by internal knowledge source: engineering, finance, and shop management. Supply chain cost modeling, engine performance, and economic scenarios are proposed categorizations that reflect three distinct types of problems which are related to different internal divisions of the company. The Model Development team will be given a problem by the Business and is capable of providing directly integrated engineering knowledge to develop a systems level view of the problem before scoping the work for an external data science partner.

The Application Support team supports both internal and external customers of Pratt & Whitney and is aligned by sub-team according to the type of customer, rather than the specific customer. This team is responsible for training, documentation, auditing and support of all data reports and models provided to the company and external customers. The Application Support team is aligned by customer group in order to train the team members toward the needs of that business sector. They may approach different segments of the model development team based upon the type of model required to address a business need. The Application Support team represents the direct monitoring of the models with the customer, and reports systematic errors back to the Model Development team for future development work.

Figure 25: Proposed Organizational Structure for Data Analytics

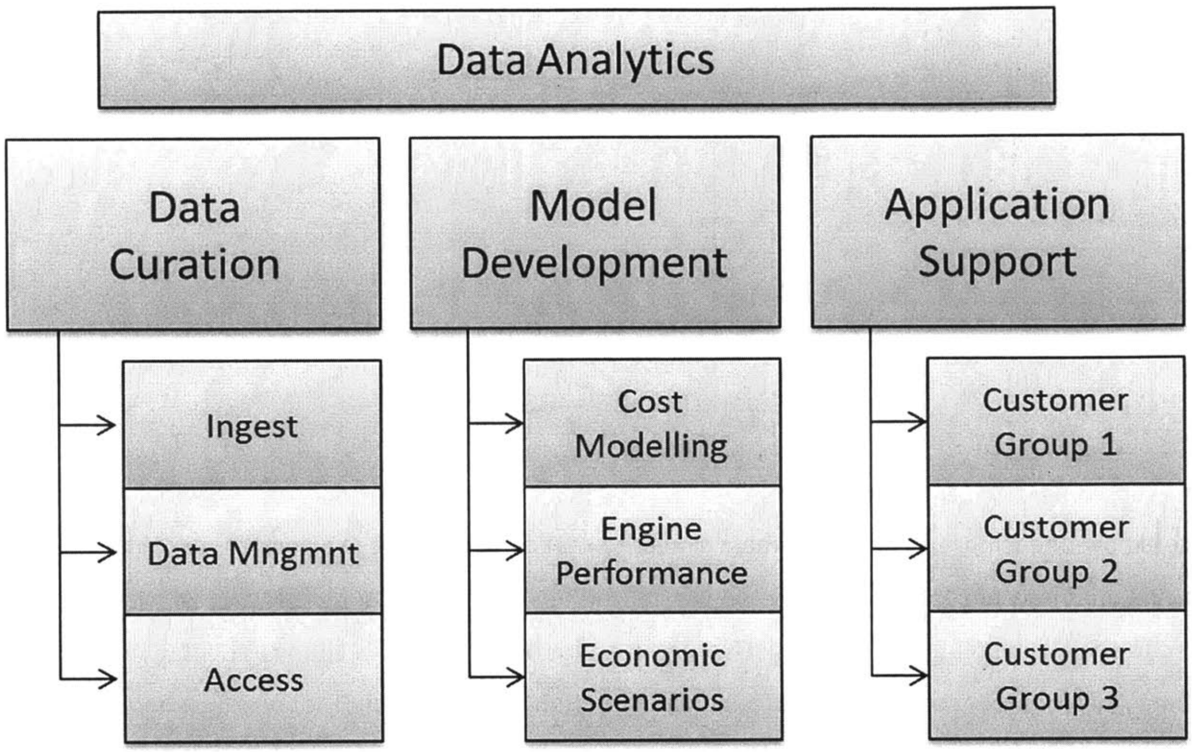

Centralization of work allows for economies of scale in the form of improvement of standard practice, and a reduction in duplicated data. The key benefits of this centralized team are found in the application of a standard method for problem solving, a central location for data warehousing and concerted effort to promulgate new tools within the company. If the effort to centralize data analytics is taken beyond these areas of synergy, it risks deterring creativity and development within the wider company. Data modeling groups that are dedicated to individual business units have unique insight to that business area, coupled with domain knowledge of the engineering specific to those problems. Integration of these teams to a central unit would be deleterious to that unique knowledge which provides their advantage in the long run. Rather the data analytics team is developed to leverage those aspects of the model development function that benefit from centralization and to provide these benefits to all other modeling groups within the company. Figure 26 shows the functional relationships of such an engagement where solid lines indicate direct communication or prime point of contact and dashed lines demonstrate oversight or occasional secondary communication.

Figure 26: Data Analytics Supporting Role to Distributed Analytics Teams

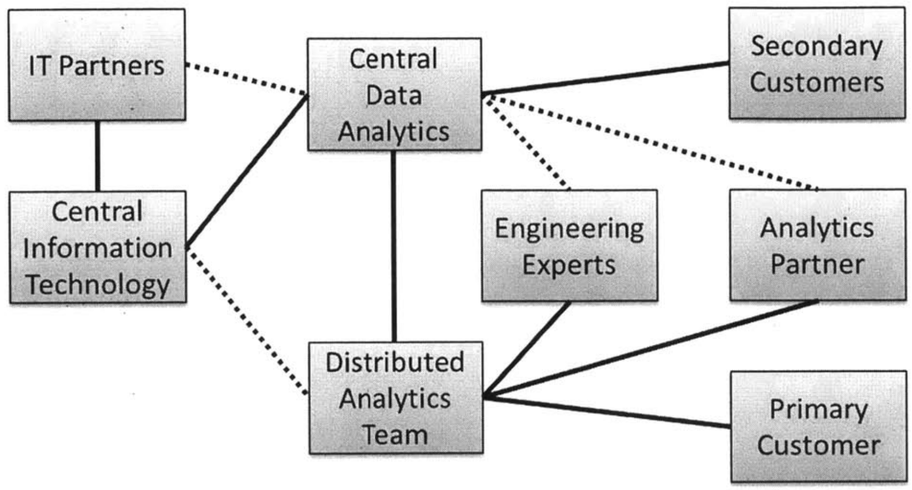

The process flow of each project would remain similar to that demonstrated in Figure 24, except that the Data Analytics lane becomes shared by the centralized team and the distributed team which is integrated to the business unit. Under this structure, the distributed analytics team relies upon the central Data Analytics team to provide thorough data curation in pairing with IT for any new data developed in the project. Advertising of the newly developed model and curated data to new customers in other business units is also provided by Data Analytics. If engineering experts or 3rd party analytics partners are needed with whom the distributed team is not familiar, the central Data Analytics team may assist in making those connections and ensuring that standard work regarding model development is applied. The distributed analytics team is itself responsible for all work with the primary model customer, as well as development and deployment of the new model in collaboration with engineering experts. The distributed analytics team provides documentation and training, and performs the continuous monitoring of the new system in collaboration with the primary customer to identify new requirements. With central support provided to each organization, new distributed teams will be easier to develop and can be funded readily by the business unit in which they reside. The team would look functionally like a single sub-team of the modeling branch of the Central Data Analytics team. The group requires perhaps three to five employees with backgrounds in engineering or business and an interest in data modeling methods. Since advanced data modeling and IT functions are performed by third parties, these employees retain skills and career options within their business units while providing unique domain knowledge to the model development process. It is important that the company become capable of handling data in both a disciplined and distributed manner. The proposed framework strives to balance these two objectives.

5.0 Conclusions

5.1 Specific Recommendations

5.1.1 Implications to Contract and Customers

In certain cases within the data, the effect of atmospheric changes between these cities observable only by satellite data and runway characteristics is equivalent to the interval effect of a 10 degree Celsius difference despite the two airports being the same temperature. The significant effect of atmospheric severity on airline operations requires a more advanced review of contract language and the customer relationship. For many years customers have been familiar with temperature, thrust, and flight length as factors for evaluating engine maintenance requirements[11]. The effects of humidity, thrust reversers, dust and elevation are not commonly regarded as negotiable contract terms. Beyond these factors, a number of others are likely to exist, and the distinction between chemical composition of aerosols is not yet considered. The addition of all necessary factors to the contracts is likely to be untenable due to the intractable nature of an N-dimensional problem. Combining these factors together in a model as shown here is capable of producing reasonably stable effects at the city level without fully understanding the damage drivers. Contract negotiations could form around lists of cities for which limited service is permitted. In these harsh environment cities new prices could be established per flight with rates updated periodically as both observed atmospheric conditions and model training evolve. Categorizing cities within the range of an airline into groups with maintenance price per city captures all current and future potential effects of the atmospheric models while enabling effective day to day decision making at the airline operations and fleet planning teams without burdensome contract restrictions on individual parameters. It is recommended that continued work be performed with contracts and legal teams to pursue a city severity pricing method for future contracts as the effect is observed to be more significant than temperature in determining airport to airport severity.

5.1.2 Corollary Information from Current Model

Continuous interval maintenance action and cleaning actions may be evaluated for impact on the engine interval using the current model. Continuous interval maintenance is included in the second model layer since it is related to the inspection policies of the airlines. Typical gains observed in interval per engine removal are 10% per removal. Within the data this varies based upon both the relative age at removal and the type of work performed. Using existing cost structures of the company, the average cost of such a visit is well established and the fleet manager is partially responsible for the joint decision with the airline about what level of work to perform on the engine when it is covered under maintenance contract. By comparing the expected revenues of a 10% life extension on the engine in question to the repair price, the fleet manager is better able to inform their decisions. Although individual engine performance varies, this information enables fleet level planning regarding the incentives placed in the contract for increased engine inspections. It is recommended that this shadow price of 10% life per visit be compared against internal processes for contract definition on the frequency and coverage of continuous interval off-wing maintenance.

The data regarding cleaning actions included in this study is present for 40% total engine intervals and uses standard practice assumptions for other airlines due to limited data availability. Although the variance inflation factor of the cleaning action coefficient is the lowest of any factor (1.7) there exists significant factor correlation with other model parameters within the data for which cleaning action was well known. The relative value of the cleaning actions most strongly correlates with A21 mid-range dust levels (0.46), Elevation Low (-.35) and Thrust Reverser usage (0.34) which have standard beta of -22.4, -21.8 and -5.8 respectively. The directionality of the cleaning action coefficient remains robust to changes in the data sampling and throughout the stepwise regression selections. Together we take this to understand that cleaning action is not detrimental to the engine life, but is performed more frequently upon engines in harsh environments. When cleaning action is removed from the regression model, the coefficients of other aerosol factors are observed to increase indicating a restorative effect of cleaning action may be masking the true effect of the aerosols. However, this shift in coefficients is not statistically significant when comparing between engines with frequent cleaning and infrequent cleaning. It is only observed between engines with cleaning when compared to engines with no cleaning at all. It is recommended that additional work be done to monitor the health of fleets before and after changes in cleaning policy within the same operating parameters. The results identified here do not justify widespread changes to the fleet cleaning policies, but support the continuation of targeted cleaning in harsh atmospheric environments.

5.1.3 The Value of Data on Future Engines

Commercial engines and airlines have an increasing number of sensors installed throughout them with each technological generation. Each sensor incurs additional cost not only for production and maintenance of the hardware but also for transmission and storage costs of data. This type of analysis helps to demonstrate which data elements provide the most value to the company. With regard to future data acquisition for engine maintenance, three areas of concern are highlighted by the proposed model: Aerosols, Humidity and Thrust Reversers.

Aerosol size, composition and density are indicated to have significant effects on engine life expectancy. In the extreme conditions, the contributions of aerosol composition city to city are found to have a difference in life by half, while within nearby cities the variance accounts for 20% difference. The aerosol information used for this work relied upon average monthly levels as determined by a limited number of satellite observations at solar noon. Acquiring improved aerosol readings for hub airports through sponsoring of ground based research may enable improved predictive maintenance. The collection of aerosol data could also be acquired by sampling the air filters from the cabin intake at each replacement. The OEM could pay for the return of these air filters for closer analysis and development of improved prediction models.

Mathematically relative humidity has been shown to have negative effects on the propulsion efficiency, and is also correlated inversely to aerosol levels in each region. The exact relationship is subject to a high degree of uncertainty. The measurement of relative humidity may be obtained in-flight by either the airframe or the engine. However, the humidity level is also a typical condition parameter recorded by the airport or local weather service on an hourly basis. The procurement and inclusion of these data at the flight by flight time level would improve the capability of the regression model to properly isolate the effects of humidity. Finally, a complete physics model of the engine at different levels of relative humidity may enable the isolation of humidity's performance effect from its air-cleansing effect. This model could be performed readily as existing engineering models contain a humidity parameter which has not been fully explored.

Thrust reverser usage was identified as a driver of engine damage. It is suspected that this damage is related to the type of aerosol and runway condition present during the landing. Improved data records of the engine performance during thrust reverser deployment may enable a better application of cross terms in the model. This data could be obtained by acquiring runway level contaminant samples through field representatives. A dataset of relative runway cleanliness could then be combined with the thrust reverser data to determine the nature of the damage. Either this is related to thrust reverser use on contaminated runways, or it is simply damage driven from the amount of time at power.

5.1.4 Team Structural Alignment

The aftermarket organization relies heavily upon data and models for decision making which requires effective development of new models and effective promulgation of data throughout the organization. IT organizations are well structured for the administration and preservation of business critical data, while individual business units are better aware of the applications of the data to create business value.

Assigning business level oversight of the Access, Data Management, and Pre-ingest functions to a centralized data analytics organization enables the advancement of data within the company while ensuring that the business side staff remain career mobile and avoid a replication of IT skills. The data analytics team outlined in Figure 25 employs the advantages of both functional alignment and customer alignment. For those aspects of the work where domain specific knowledge such as programming or advanced modeling methods are required, the data analytics team relies upon external vendors. Meanwhile, the skills cultivated within the company's aftermarket business unit remain consistent with the core engineering skills of the OEM and center on foundational knowledge of the product. The data analytics team provides a centralized gateway from any business unit into an array of data analytics resources both inside and outside the company. Individual teams throughout the aftermarket organization will be empowered to continue developing queries, monitoring data, and creating and using models.

5.2 General Implications to World

5.2.1 Application of Airport Severity to Airline Industry

Airline maintenance and overhaul costs for airframe and engine are typically observed at 10-12% of the airline total cost structure[56]. Of these costs, the engines are nearly half, which places the cost of engine maintenance and overhaul on an equivalent volume to the current profit margin of successful airlines. This research has shown that beyond the industry established effects (temperature, thrust and flight length), atmospheric effects, inspection policy and thrust reverser policies are responsible for between 20% and 50% differences in similar fleet interval. If we assume that overhaul costs are equivalent regardless of TOW, this indicates that the cost of these unknown factors is between 1-3% of total airline cost structure. This is equivalent to the typical revenue fluctuation caused by international currency management (1%), or shifts in the annual storm season (1.5%) in 2014 for Delta Airlines[81]. The fluctuation in revenues between flights on peak travel hours and off-peak travel hours significantly outweighs the maintenance costs of taking off in higher temperatures[56]. However, when cost structures for competing airports in a region are considered, it may be beneficial to include atmospheric costs in the fleet planning decision. For example, when two airports in a popular Mexican tourist destination separated by less than 20 miles are tested in the proposed model, one airport has half the expected engine life of the other. On a typical mission this would incur a premium cost of approximately \$1000 per round trip or approximately \$150 per block hour which would contribute to fleet optimization decisions.

5.2.2 Implications of Aerosol Effects and Incentives

Airports generate revenue from vendors and airlines. The airline revenue is acquired by typically by a landing fee structure or per passenger fee or combination of the two, with average total cost per enplanement in the US ranging from $2.80 at Charlotte Douglas to $24.20 at New York JFK. Studies have observed that the prices are typically set by the open market as a function of airport size, regional demand and competition[82]. We find that the cost of engine maintenance incurred by severity per enplanement is similar in price and range to the range of costs in US airports. Figure 27 shows the spread of maintenance costs per enplanement derived from the current model with an assumed a single-isle aircraft for typical mission of 85% load factor. If this cost structure were committed to contracts by the OEM or incurred directly by airlines which pay for their own maintenance expenses, then the airport demand situation would shift. Those airports whose geography naturally incurs higher costs on the airline would face additional negative pressure on their fee structure which responds to market forcing particularly in secondary hubs and regional airports. The standard deviation in Figure 27 is 70¢ or 25% of the typical fee for Charlotte Douglas. Larger airports like JFK could be expected to disregard this effect due to their strong market position. Smaller airports such as Charlotte Douglas may face revenue erosion if their severity is found to be significant.

Figure 27: Engine Maintenance Cost Per Enplanement of Studied Cities in the Current Model

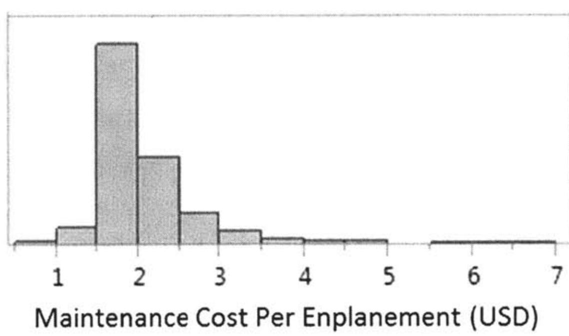

5.2.3 Applications of Multiple Information Source to Other Industries

A multiple information source system employs a variety of modeling techniques and decision gates in order to enable effective representation of highly complex systems. While individual modeling approaches risk oversimplification of the underlying interactions of data, multiple information source applies the most applicable model to only that part of the data for which the model type applies. This enables decomposition of complex socio-technical systems into an array of distinct solvable problems. These methods have been widely adopted in long range economic and development forecasts and more recently through the insurance industry to health care deployment. This work has shown the application of multiple information source to a unique engineering problem overlaid with both socio-political factors and unstructured environmental effects. One key advantage of multiple information source hierarchical models over complex single form models entails the ability of the model to be communicated effectively to 3rd party customers or auditors. As data systems expand and modeling complexity increases, multiple information source systems will enable the development of data systems that are both complex and relatable. Individual layers of the system may be interchanged with newer models without compromising the entire system, though individual layers may require re-training.

5.3 Future Work

5.3.1 Extension of Data To Other Business Problems

Three topics for future work take advantage of adjacency to the problem solved here. First, the total cost of repair for any product could be related by this parametric method to its operating environment. This relationship is likely to be complex and parametric in nature. Correlation studies between repair costs and environmental effects may yield additional insights. This work is important to the application of any survival time model as it is possible that effects which drive early failure in a pipe or appliance may be causing targeted damage that is easier to repair, or that they may be causing greater systemic damage that is more expensive to repair.

Second, the relationship between product origin and product lifetime could be explored. The environmental damage effects of geographically diverse operations are often overlapping with geographic disparity in manufacturing quality. This analysis may enable a better isolation of product quality drivers based upon the environment that any product is expected to operate in. Although much of this may already be optimized by tacit knowledge within a global company, financial value of off-shore vs. on-shore production regularly depends upon accurate estimates of production costs and production quality.

Third, this enables better risk management of any fleet. This could cover a trucking fleet, airline operation or military applications with large numbers of similar products facing diverse environmental hazards. Although the results of the current model are aggregated by fleet for analysis of fleet level risk, the output of the system can be applied at an engine, or single entity, level to generate improved maintenance forecasts by region in the coming months. This type of analysis could lead to improved visit forecasts, and if linked to the cost analysis above, could enable proactive resource planning.

5.3.2 Defining New Data Acquisition Targets

The ability exists within ACARS message architecture for additional data to be transmitted. The value of new data and additional information about existing data streams is not always clear. This can lead to a tendency to either hoard all data possible at great expense, or disregard data that does not pose direct value to the organization. An appropriate middle ground between these can be obtained by structured analysis of the system. In the instance of engine performance, the current model demonstrates the value of certain data and points to the value of unknown data. For example, improved data granularity with respect to climb profiles and performance may reduce the assumptions used in the model. At the same time it does not appear necessary to track every flight parameter continuously in order to determine engine health and life expectancy. This is shown by the fact that the current model does not improve significantly when all performance characteristics such as vibration and module performance ratios are included in the regression. As future event prediction models are developed, a system dynamics view will help identify missing variables and inform the prioritization of changes to ACARS data. These models should also be challenged once created to perform analysis using less data than is currently available. In this way the model sensitivity to missing data can be established and the actual value of data collection can be better known by the company.

Appendix

NASA Satellite Data Acquisition

The NASA satellite TERRA (EOS AM-1) performs a polar circular low earth orbit with a sun-synchronous repeating ground track with earth facing cameras measuring solar reflectance of the atmosphere and surface. Two modules from TERRA provided data to this research: Moderate-Resolution Imaging Spectroradiometer (MODIS) and Multi-angle Imaging SpectroRadiometer (MISR). Data from MODIS is used during on board calculations in MISR, and off board calculations by NASA in preparation of MISR datasets. MODIS is responsible for the identification of ground terrain types which are provided to MISR as a baseline for calculated expected solar reflection. In addition, MISR relies upon atmospheric and surface climatology (TASC) and Radiometric Camera-by-Camera CloudMask Thresholds (RCCT) to determine baseline atmospheric values such as temperature, wind speed and the presence and height of clouds. When clouds fully obscure the pixel and prevent observation of the low atmospheric aerosols, MISR optical depth calculations are omitted. MISR uses four cameras and the time correlated datasets to compute both the Aerosol Optical Depth (AOD) and water column present in each pixel. The pixel level values are recorded in data level 3 products on the NASA server. MISR aggregates values across an area approximately 1.1 km by 1.1 km to perform mixture analysis. The level 3 readings of AOD across 4 wavelengths observed for each pixel could be caused by a variety of aerosol compositions. The satellite is programmed by NASA to check the AOD against models using Mie Theory equations for a set of assumed aerosol composition models. The composition with best fit to the data across the 1.1 km square region is provided to the level 2 dataset MIL2ASAE. For the determination of aerosol types in each observation period, this research selected the file RegLowestResidMixture from the MIL2ASAE product. Once the mixture has been selected, the AOD calculations at each of the four wavelengths provided by MISR will exhibit a degree of error from one another.

For the purposes of data self-consistency, all atmospheric observations in this research were made using the camera observations at wavelength 558 nm since it is closest to solar peak irradiance and therefore subject to the least overall error and is given the highest weighting by NASA's calculations[83]. For each aerosol in the mixture composition models, NASA's dataset provides the specific effective radius (r_{eff}) and cross-sectional absorption coefficient (σ). For this research, data from the MIL2ASAE dataset was processed using Matlab 2014a The code translates readings into column average volume mass loading of aerosols within a 30 nautical mile radius of each airport and was developed as a part of this study. Aerosol optical depth is translated into column volume mass loading (VL) by the following equation:

Equation 8 $VL_i = \frac{4}{3}\pi * {r_{eff}}_i^{\,3} * \frac{AOD}{\sigma_i} * f_i$

where f_i represents the volume based fraction for the individual aerosol in the composition model of best fit, and VL is calculated for each aerosol present in the mixture. Although MISR has the functionality for 21 aerosol definitions, a total of 8 aerosol definitions were in use by the NASA models at the time of this research as given in Table 10. Volume Mass Loading may be extended to the more conventional Mass Loading by the particle density. This conversion of units is an unnecessary step since the nature of the current study is to assign a variable coefficient to each term and therefore further in regard to aerosol effects is based on the use of Column Average Volume Mass Loading in units of μm.

Table 10: Aerosol types modeled by NASA MISR

Aerosol Code	Type	Size (μm)	Expected Composition
1	nonabsorbing	0.06	sulfate/organic
2	nonabsorbing	0.12	sulfate/organic
3	nonabsorbing	0.26	sulfate/organic
6	nonabsorbing	2.8	salt/organic
8	absorbing	0.12	sulfate/organic (ssagreen.9)
14	absorbing	0.12	sulfate/organic (ssagreen.8)
19	grains	0.75	mode1 dust
21	spheroidal	2.4	mode2 dust

Relative humidity and dew point are also derived from the MIL2ASAE dataset using the atmospheric reference data from TASC. All calculations are made at the 1.1 km square level provided by MIL2ASAE dataset and then aggregated further to the airport level to a single set of values for each airport. The central latitude and longitude of each airport is acquired and used to filter the MIL2ASAE dataset[84]. All observations are grouped by orbit and airport to yield the average volume mass loading per aerosol type within a 30 nm radius. Standard commercial aircraft approach airports through the top of the Class B airspace which has a height of 10,000 feet and a maximum radius of 30 nm[85]. NASA MISR calculations assume full extinction of measured aerosols above 10,000 feet for all categories; therefore a 30 nm radius is established around each airport as the data collection region.

The ground track of TERRA repeats every 233 orbits or approximately 16 days, with an orbital period of 98.8 minutes. The ground track area is designed for minimal equatorial overlap, so that most equatorial cities are observed in either a single track or two tracks providing atmospheric data for equatorial airports every 16 days or more frequently in some cases. Polar locations are observed by every track, and cities above tropic zones are commonly observed in more than one of the orbital paths.

Correlation Matrix of Primary Factors

	Flight Length	% Derate	ISA Adj Temp	Elevation	Cruise Alt.	Takeoff Weight	Runway Length	TR_Usage	Rel. Humidity	A1	A2	A3	A6	A8	A14	A19	A21
Flight Length	1.00	-0.16	-0.12	-0.34	0.08	0.45	-0.02	0.30	-0.01	-0.07	-0.36	-0.37	-0.47	-0.26	0.05	-0.22	0.31
% Derate	-0.16	1.00	-0.07	-0.02	0.23	-0.81	-0.05	-0.26	-0.05	-0.05	-0.12	-0.12	-0.15	-0.06	-0.01	-0.17	-0.09
ISA Adj Temp	-0.12	-0.07	1.00	0.33	0.01	-0.01	-0.07	0.08	0.54	0.44	0.12	0.19	0.25	0.15	0.29	0.13	0.09
Elevation	-0.34	-0.02	0.33	1.00	0.25	-0.41	0.27	-0.02	-0.19	0.32	-0.17	-0.10	0.05	-0.19	-0.09	-0.28	-0.16
Cruise Alt.	0.08	0.23	0.01	0.25	1.00	-0.24	-0.03	0.04	-0.04	0.01	-0.58	-0.64	-0.63	-0.26	0.21	-0.48	0.36
Takeoff Weight	0.45	-0.81	-0.01	-0.41	-0.24	1.00	-0.15	0.49	0.28	-0.22	0.15	0.10	0.03	0.17	0.16	0.30	0.31
Runway Length	-0.02	-0.05	-0.07	0.27	-0.03	-0.15	1.00	-0.50	-0.25	0.32	0.02	0.03	0.08	-0.04	-0.31	-0.05	-0.12
TR_Usage	0.30	-0.26	0.08	-0.02	0.04	0.49	-0.50	1.00	0.37	-0.45	0.05	-0.02	-0.15	0.08	0.36	0.16	0.37
Rel. Humidity	-0.01	-0.05	0.54	-0.19	-0.04	0.28	-0.25	0.37	1.00	-0.29	0.46	0.38	0.21	0.55	0.60	0.53	0.47
A1	-0.07	-0.05	0.44	0.32	0.01	-0.22	0.32	-0.45	-0.29	1.00	-0.25	-0.10	0.13	-0.31	-0.25	-0.35	-0.38
A2	-0.36	-0.12	0.12	-0.17	-0.58	0.15	0.02	0.05	0.46	-0.25	1.00	0.84	0.78	0.73	0.11	0.89	-0.03
A3	-0.37	-0.12	0.19	-0.10	-0.64	0.10	0.03	-0.02	0.38	-0.10	0.84	1.00	0.88	0.53	-0.07	0.76	-0.20
A6	-0.47	-0.15	0.25	0.05	-0.63	0.03	0.08	-0.15	0.21	0.13	0.78	0.88	1.00	0.50	-0.14	0.68	-0.34
A8	-0.26	-0.06	0.15	-0.19	-0.26	0.17	-0.04	0.08	0.55	-0.31	0.73	0.53	0.50	1.00	0.38	0.68	0.11
A14	0.05	-0.01	0.29	-0.09	0.21	0.16	-0.31	0.36	0.60	-0.25	0.11	-0.07	-0.14	0.38	1.00	0.20	0.38
A19	-0.22	-0.17	0.13	-0.28	-0.48	0.30	-0.05	0.16	0.53	-0.35	0.89	0.76	0.68	0.68	0.20	1.00	0.26
A21	0.31	-0.09	0.09	-0.16	0.36	0.31	-0.12	0.37	0.47	-0.38	-0.03	-0.20	-0.34	0.11	0.38	0.26	1.00

References

[1] E. Okamoto and E. Hata, "[Estimation of disease-specific costs in a dataset of health insurance claims and its validation using simulation data]," *Nihon Kōshū Eisei Zasshi Jpn. J. Public Health*, vol. 50, no. 12, pp. 1135–1143, Dec. 2003.

[2] P. Muennig, P. Franks, and M. Gold, "Research articles: The cost effectiveness of health insurance," *Am. J. Prev. Med.*, vol. 28, pp. 59–64, Jan. 2005.

[3] T. L. Kolbe-Alexander, J. Conradie, and E. V. Lambert, "Clustering of risk factors for non-communicable disease and healthcare expenditure in employees with private health insurance presenting for health risk appraisal: a cross-sectional study," *BMC Public Health*, vol. 13, pp. 1213–1213, Dec. 2013.

[4] J. C. J. H. Aerts and W. J. W. Botzen, "Climate change impacts on pricing long-term flood insurance: A comprehensive study for the Netherlands," *Glob. Environ. Change*, vol. 21, pp. 1045–1060, Jan. 2011.

[5] Chin-Piao Yeh, Ai-Chi Hsu, Wei-Hsien, and Kuang-Cheng Chai, "Neural Network Forecasts of Taiwan Bureau of National Health Insurance Expenditures," *Int. J. Bus. Finance Res. IJBFR*, vol. 8, no. 5, pp. 95–114, Jun. 2014.

[6] E. Okamoto and E. Hata, "Estimation of disease-specific costs in health insurance claims: a comparison of three methods," *Nihon Kōshū Eisei Zasshi Jpn. J. Public Health*, vol. 51, no. 11, pp. 926–937, Nov. 2004.

[7] S. Tanihara, Z. Yamagata, and H. Une, "[Reliability of health insurance claim statistical data based on the principal diagnosis method]," *Nihon Eiseigaku Zasshi Jpn. J. Hyg.*, vol. 63, no. 1, pp. 29–35, Jan. 2008.

[8] C. Markou and G. Cros, "Airline Maintenance Cost Executive Commentary (FY2011)," International Air Transport Association, Apr. 2013.

[9] Rolls-Royce, "Rolls-Royce celebrates 50th anniversary of Power-by-the-Hour," *Rolls-Royce News*, 30-Oct-2012. [Online]. Available: http://www.rolls-royce.com/news/press_releases/2012/121030_the_Hour.jsp. [Accessed: 21-Sep-2014].
[10] K. M. Ravi and K. H. Rajesh, "Reliability Analysis," QuEST Global Services PTE LTD, Singapore, White Paper, 2014.
[11] S. Ackert, "Engine Maintenance Concepts for Financiers," Jackson Square Aviation, 2nd, Sep. 2011.
[12] The World Bank and I. C. A. O. ICAO, "Air Transport, Passengers Carried," *Data*, 2014. [Online]. Available: http://data.worldbank.org/indicator/IS.AIR.PSGR?page=1. [Accessed: 21-Sep-2014].
[13] B. Millhaem, "Engine operation in harsh environments," presented at the MRO MIDDLE EAST 2011, Dubai, UAE, 02-Feb-2011.
[14] "01_AVM_Cover1_Dec14Jan15_Signoff.indd - AVM-2014-12-DIGITAL.pdf." [Online]. Available: http://www.avm-mag.com/wp-content/mag/PDFs/AVM-2014-12-DIGITAL.pdf. [Accessed: 31-Jan-2015].
[15] cjh@dongyang. ac. k. Jinhyung Cho1, "Credibility Enhancement of Social Curation Using Multiple Information Sources," *Int. J. Adv. Comput. Technol.*, vol. 5, no. 13, pp. 409–414, Sep. 2013.
[16] E. Hato, M. Taniguchi, Y. Sugie, M. Kuwahara, and H. Morita, "Incorporating an information acquisition process into a route choice model with multiple information sources," *Transp. Res. Part C Emerg. Technol.*, vol. 7, no. 2–3, pp. 109–129, Apr. 1999.
[17] Y. Y. Tang, Y. Z. Qu, and C. Y. Suen, "Multiple-level information source and entropy-reduced transformation models," *Pattern Recognit.*, vol. 24, no. 4, pp. 341–357, 1991.
[18] E. Ramasso and R. Gouriveau, "Remaining Useful Life Estimation by Classification of Predictions Based on a Neuro-Fuzzy System and Theory of Belief Functions," *IEEE Trans. Reliab.*, vol. 63, no. 2, pp. 555–566, Jun. 2014.
[19] R. Moghaddass and M. J. Zuo, "An integrated framework for online diagnostic and prognostic health monitoring using a multistate deterioration process," *Reliab. Eng. Syst. Saf.*, vol. 124, pp. 92–104, Apr. 2014.
[20] S. Somasundaram and R. Krishnan, "The Reliability Analysis of a 5-Component Dynamic Warm Standby System with Sensor and Switch," *Int. J. Stat. Syst.*, vol. 4, no. 1, pp. 13–23, Mar. 2009.
[21] Z. Dai and Z. Wang, "Protection Dynamic Reliability Analysis System Based on 3RF Technique," *IEEE Trans. Power Syst.*, vol. 26, no. 3, pp. 1137–1144, Aug. 2011.
[22] G.-Y. Chung, J. V. R. Prasad, M. Dhingra, and R. Meisner, "Real Time Analytical Linearization of Turbofan Engine Model," *J. Eng. Gas Turbines Power*, vol. 136, no. 1, pp. 1–13, Jan. 2014.
[23] B. Milstein, J. Homer, and G. Hirsch, "Analyzing National Health Reform Strategies With a Dynamic Simulation Model," *Am. J. Public Health*, vol. 100, no. 5, pp. 811–819, May 2010.
[24] T. G. Schmitt, M. Thomas, and N. Ettrich, "Assessment of urban flooding by dual drainage simulation model RisUrSim," *Water Sci. Technol.*, vol. 52, no. 5, pp. 257–264, Sep. 2005.
[25] "4: Gas Turbine Major Components and Modules," *Gas Turbines*, pp. 119–166, Jan. 2007.
[26] F. Brooks, "GE Gas Turbine Performance Characteristics," *GE Power Syst.*, no. GER-3567H.
[27] Rolls Royce, *The Jet Engine*, 5th ed. Renault Printing Co Ltd, 1986.
[28] Pratt & Whitney, *The Aircraft Gas Turbine Engine and Its Operation*. Pratt & Whitney Aircraft, 1974.
[29] R. A. Roberts and S. M. Eastbourn, "Modeling Techniques for a Computational Efficient Dynamic Turbofan Engine Model," *Int. J. Aerosp. Eng.*, pp. 1–11, Jan. 2014.
[30] D. L. Simon and J. B. Armstrong, "An Integrated Approach for Aircraft Engine Performance Estimation and Fault Diagnostics," *J. Eng. Gas Turbines Power*, vol. 135, no. 7, pp. 071203–071203, Jun. 2013.
[31] J. S. Litt, "An Optimal Orthogonal Decomposition Method for Kalman Filter-Based Turbofan Engine Thrust Estimation," *J. Eng. Gas Turbines Power*, vol. 130, no. 1, pp. 1–12, Jan. 2008.

[32] D. Simon and D. L. Simon, "Kalman filtering with inequality constraints for turbofan engine health estimation," *IEE Proc. -- Control Theory Appl.*, vol. 153, no. 3, pp. 371–378, May 2006.
[33] Feng Lu, Jinquan Huang, and Yiqiu Lv, "Gas Path Health Monitoring for a Turbofan Engine Based on a Nonlinear Filtering Approach," *Energ. 19961073*, vol. 6, no. 1, pp. 492–513, Jan. 2013.
[34] H. Hanumanthan, A. Stitt, P. Laskaridis, and R. Singh, "Severity estimation and effect of operational parameters for civil aircraft jet engines," *Proc. Inst. Mech. Eng. -- Part G -- J. Aerosp. Eng. Sage Publ. Ltd*, vol. 226, no. 12, p. 1544, Dec. 2012.
[35] "CFM56-3 maintenance analysis & budget," *Aircarft Commer.*, no. 45, pp. 18–28, Apr. 2006.
[36] "2014 Breakthrough: The Market Changing Pratt & Whitney Geared Turbofan Engine," *AirInsight.* [Online]. Available: http://airinsight.com/product/breakthrough-market-changing-pratt-whitney-geared-turbofan-engine/. [Accessed: 09-Feb-2015].
[37] R. Mishra and C. Beura, "Life Consumption Assessment of a Large Jet Engine," *J. Fail. Anal. Prev.*, vol. 14, no. 4, p. 519, Aug. 2014.
[38] M. Johnston, C. Campbell, R. Hayward, M. Lowerison, V. K. Noonan, T. Pfister, C. Maxwell, C. M. Fortin, E. E. Smith, J. K. Mah, M. K. Kapral, N. Jette, T. Pringsheim, and L. Korngut, "Registry data storage and curation," *Can. J. Neurol. Sci. J. Can. Sci. Neurol.*, vol. 40, no. 4 Suppl 2, pp. S35–S40, Jul. 2013.
[39] P. Lord, A. Macdonald, L. Lyon, and D. Giaretta, "From Data Deluge to Data Curation." UKOLN, University of Bath.
[40] "Recommendation for Space Data System Practices," CCSDS, Recommended Practice 650.0-M-2, Jun. 2012.
[41] D. Sawyer, "The Open Archival Information System and the NSSDC." NASA.
[42] A. Palaiologk, A. Economides, H. Tjalsma, and L. Sesink, "An activity-based costing model for long-term preservation and dissemination of digital research data: the case of DANS," *Int. J. Digit. Libr.*, vol. 12, no. 4, pp. 195–214, Sep. 2012.
[43] P. Laughton and T. du Plessis, "Data Curation in the World Data System: Proposed Framework," *Data Sci. J.*, vol. 12, p. 56, Dec. 2013.
[44] P. Laughton, "OAIS functional model conformance test: a proposed measurement," *Program*, vol. 46, no. 3, pp. 308–320, 2012.
[45] P. M. Long and R. A. Servedio, "Random classification noise defeats all convex potential boosters," *Mach. Learn.*, vol. 78, no. 3, pp. 287–304, Mar. 2010.
[46] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, Jun. 2006.
[47] G. Reinman, "Pratt & Whitney Interview," Aug-2014.
[48] X. Liu, R. C. Easter, S. J. Ghan, R. Zaveri, P. Rasch, X. Shi, J.-F. Lamarque, A. Gettelman, H. Morrison, F. Vitt, A. Conley, S. Park, R. Neale, C. Hannay, A. M. L. Ekman, P. Hess, N. Mahowald, W. Collins, M. J. Iacono, and C. S. Bretherton, "Toward a minimal representation of aerosols in climate models: description and evaluation in the Community Atmosphere Model CAM5," *Geosci. Model Dev.*, vol. 5, no. 3, pp. 709–739, Jul. 2012.
[49] S. Bennett, "Log-Logistic Regression Models for Survival Data," *J. R. Stat. Soc. Ser. C Appl. Stat.*, vol. 32, no. 2, pp. 165–171, Jan. 1983.
[50] A. Cetin and A. Naess, "Toward a proper statistical description of defects," *Int. J. Fatigue*, vol. 38, pp. 100–107, May 2012.
[51] "The power behind the best decisions in global aviation," *Official Airline Guide - Absolute Aviation Advantage*. [Online]. Available: http://www.oag.com/. [Accessed: 06-Feb-2015].
[52] E. L. Kaplan and P. Meier, "Nonparametric Estimation from Incomplete Observations," *J. Am. Stat. Assoc.*, vol. 53, no. 282, pp. 457–481, Jun. 1958.

[53] G. E. P. Box and N. R. Draper, "Design Aspects of Variance, Bias, and Lack of Fit," in *Response Surfaces, Mixtures, and Ridge Analyses*, John Wiley & Sons, Inc., 2007, pp. 413–459.
[54] M. Gell, D. N. Duhl, and A. F. Giamei, "The Development of Single Crystal Superalloy Turbine Blades - 01-1022-205.pdf." Pratt & Whitney Aircraft, 1980.
[55] H. Jiang and C. Barnhart, "Robust airline schedule design in a dynamic scheduling environment," *Comput. Oper. Res.*, vol. 40, no. 3, pp. 831–840, Mar. 2013.
[56] P. Belobaba, A. Odoni, and C. Barnhart, *The Global Airline Industry*, 1st ed. John Wiley & Sons, Ltd, 2009.
[57] M. Pidwirny, "Global surface temperature distribution," *The Encyclopedia of Earth*, Oct-2007. [Online]. Available: http://www.eoearth.org/view/article/51cbede47896bb431f694887/. [Accessed: 12-Feb-2015].
[58] O. Altaratz, I. Koren, L. A. Remer, and E. Hirsch, "Review: Cloud invigoration by aerosols—Coupling between microphysics and dynamics," *Atmospheric Res.*, vol. 140–141, pp. 38–60, Apr. 2014.
[59] T. Wagner, S. Beirle, D. Loyola, K. Mies, and S. Slijkhuis, "Retrieval of total precipitable water from GOME-1, SCIAMACHY and GOME-2," *The ESA DUE GlobVaport Project*, Jan-2011. [Online]. Available: http://www.globvapour.info/newsarchive.html. [Accessed: 12-Feb-2015].
[60] NASA, "NASA - New Map Offers a Global View of Health-Sapping Air Pollution," Sep-2010. [Online]. Available: http://www.nasa.gov/topics/earth/features/health-sapping.html. [Accessed: 12-Feb-2015].
[61] R. Lugo-Leyte, J. M. Zamora-Mata, M. Toledo-Velázquez, M. Salazar-Pereyra, and A. Torres-Aldaco, "Methodology to determine the appropriate amount of excess air for the operation of a gas turbine in a wet environment," *Energy*, vol. 35, no. 2, pp. 550–555, Feb. 2010.
[62] N. Y. Moore, C. A. Grammich, and R. W. Bickel, *Developing Tailored Supply Strategies*. Rand Corporation, 2007.
[63] "Discipline And Imagination: How United Technologies Chairman Louis Chenevert Keeps His Powerhouse Humming," *Forbes*. [Online]. Available: http://www.forbes.com/sites/lorenthompson/2014/04/02/discipline-imagination-how-united-technologies-chairman-louis-chenevert-keeps-his-powerhouse-humming/. [Accessed: 14-Feb-2015].
[64] S. Broderick, "Trends Suggest Solid Aftermarket Prospects For Pratt & Whitney," *Aviation Week*, 13-Jun-2014.
[65] UTC, "United Technologies Corporation 2013 Annual Report," Annual Report, Apr. 2014.
[66] C. Adams, "Making the Most of Big Data for MRO," *Aviation Maintenance*, pp. 23–28, Jan-2015.
[67] "Exhibitor`s Information: Osys." [Online]. Available: http://www.aircraft-commerce.com/conferences/MIA/VENDORS/OSyS.asp. [Accessed: 14-Feb-2015].
[68] "In-Flight," *Pratt & Whitney Engine Health Management*. [Online]. Available: http://www.pw.utc.com/In_Flight. [Accessed: 14-Feb-2015].
[69] H. K. Bowen and C. Purrington, "Pratt & Whitney: Engineering Standard Work," Harvard Business School, Pratt & Whitney, Case Study 9-604-084, Mar. 2006.
[70] A. Levin, "Airways in USA are the safest ever," *USA Today*, 30-Jun-2006.
[71] B. Cleveland, "Staff to Supervisor Ratio." International Customer Management Institue, 13-Jun-2012.
[72] AIAA, "Recruiting, Retaining, and Developing a World-Class Aerospace Workforce." AIAA, 2009.
[73] C. B. Sanders, M. D. Steward, and S. Bridges, "Facilitating knowledge transfer during SOX-mandated audit partner rotation," *Bus. Horiz.*, vol. 52, no. 6, pp. 573–582, Nov. 2009.
[74] K. Girard, "CSC, Pratt & Whitney negotiate services deal," *Technology News - CNET News*, Feb-1999. [Online]. Available: http://www.cnet.com/news/. [Accessed: 01-May-2015].
[75] CSC, "CSC Signs Major Five-Year IT Services Contract Extension With UTC," *Computer Sciences Corporation*, 18-Mar-2010. [Online]. Available:

http://www.csc.com/newsroom/press_releases/43082-csc_signs_major_five_year_it_services_contract_extension_with_utc. [Accessed: 01-May-2015].
[76] D. Pickett and J. Rice, "Pratt & Whitney Interview," Nov-2014.
[77] B. Weimar, "Assumptions about University-Industry Relationships in Continuing Professional Education: a re-assessment," *Eur. J. Educ.*, vol. 27, no. 4, p. 379, Dec. 1992.
[78] D. Smith, "Pratt & Whitney Interview," Dec-2014.
[79] A. Grace, "Pratt & Whitney Interview," Aug-2014.
[80] A. Chakravarty, A. Kumar, and R. Grewal, "Customer Orientation Structure for Internet-Based Business-to-Business Platform Firms," *J. Mark.*, vol. 78, no. 5, pp. 1–23, Sep. 2014.
[81] A. Levine-Weinberg, "Delta Air Lines, Inc.: The Dip Is a Buying Opportunity (DAL)," *The Motley Fool*, 17-Feb-2015. [Online]. Available: http://www.fool.com/investing/general/2015/02/17/delta-air-lines-inc-the-dip-is-a-buying-opportunit.aspx. [Accessed: 17-Feb-2015].
[82] Y. Y. Choo, "Factors affecting aeronautical charges at major US airports," *Transp. Res. Part Policy Pract.*, vol. 62, pp. 54–62, Apr. 2014.
[83] "MISR: Spectral Coverage." [Online]. Available: https://www-misr.jpl.nasa.gov/Mission/misrInstrument/spectralCoverage/. [Accessed: 30-Jan-2015].
[84] "Data downloads @ OurAirports." [Online]. Available: http://ourairports.com/data/. [Accessed: 30-Jan-2015].
[85] "Chapter 15. Class B Airspace." [Online]. Available: https://www.faa.gov/air_traffic/publications/atpubs/AIR/air1502.html. [Accessed: 30-Jan-2015].
[86] M. Naeem, "Implications of day temperature variation for an aero-engine's HP turbine-blade's creep life-consumption," Aerosp. Sci. Technol., vol. 13, pp. 27-35, Jan. 2009.