MANAGERIAL FACTORS AFFECTING AIRCRAFT MAINTENANCE: AN AGENT BASED MODEL AND OPTIMIZATION WITH SIMULATED ANNEALING

by

Douglas Hamilton

BA Mathematics
University of Illinois Springfield 2012

Submitted to the Faculty in partial fulfillment of the requirements for the degree of

Master of Science in Engineering and Management

at

Massachusetts Institute of Technology

January 14, 2015 [February 2015]

Author: .... Signature redacted

Douglas Maxwell Hamilton
System Design and Management Program

Certified by: .... Signature redacted

Daniel D Frey
Thesis Supervisor
Professor of Mechanical Engineering

Accepted by: .... Signature redacted

Patrick Hale
Chairman & Senior Lecturer, Engineering Systems Division
Director, System Design and Management Program

MANAGERIAL FACTORS AFFECTING AIRCRAFT MAINTENANCE: AN AGENT BASED MODEL AND OPTIMIZATION WITH SIMULATED ANNEALING

By

Douglas M Hamilton

Submitted to the System Design and Management Program on January, 15, 2015 in Partial fulfillment of the Requirements for the Degree of Master of Science in Engineering and Management.

Abstract

A single objective agent based model of managerial factors affecting aircraft maintenance was build based on a case study done regarding safety climate's effect on maintenance efficacy in the Korean Air Force. In particular the model measures the effect of managerial context and command on agents' motivation and efficacy. The model is then optimized using a simulated annealing algorithm. Input parameters were varied to ensure reliability and repeatability of results. The model's sensitivity, in terms of optimal input vector and results, were also tested across a variety of input parameters. Results suggested that across all input parameters two managerial contexts dominated: contingent reward systems and laissez-faire.

Thesis Supervisor: Daniel D Frey
Title: Professor of Mechanical Engineering

Acknowledgments

I would like to thank the following people for their contributions to both this thesis and my graduate experience at large:

- Dan Frey for advising and helping me navigate the thesis process
- The SDM front office staff for their invaluable contributions making everything (from financing to logistics) work
- Pat Hale for guidance in the early stages of finding an advisor and topic
- My family for their relentless goading to finish,
- My cohort-mates for their support in and out of the classroom

Contents

Table of Figures

Table of Equations

CHAPTER 1: Introduction

In fiscal year 2009 commercial airlines reported that, depending on aircraft size, per-flying-hour maintenance cost ran between $461 and $1430. In house labor and material costs each accounted for 21% of the total, to comprise 42% of all maintenance spending across commercial fleets. [1] From the period of 1974 to 2005 the cost of military aircraft has outpaced inflation by at least 2.5%. For certain families of airframes and measures of inflation the offset was as high as 8%. [2] The rising costs of replacing aircraft and high flying hour maintenance cost make optimizing maintenance throughput in terms of both volume and efficiency critical to a cost conscious airborne service. Improvements in maintenance outcomes can be gained by improving supply chains, reducing fatigue and stress of maintainers, or improving individual maintainer throughput. The latter two of these factors are closely related to maintence units' safety climate. [3] Safety climate is itself a function of management's commitment to safety and both the level of concern managers show for subordinate welfare and the methods by which they voice that concern. In order to understand safety climate's impact on maintence outcomes this paper will examine the relation between safety climate and outcomes, as well as the factors which improve safety climate with the goal of optimizing managerial inputs to maximize maintenance efficacy.

In order to optimize input, first a computational model of the scenario must be made. Building on the work of Park, Kang and Son this paper will present a model mapping safety climate to unit performance. It will also build a mapping from managerial context and degree of centralization in organizational decision making to safety climate. This input will be tempered by an agent based model emulating individual maintainers' responses to centralization and managerial style. The model will also be tunable across a number of other parameters relevant to organizations, but which are not easily changed by management. Since the nature of the input variables does not make the problem amenable to gradient based optimization methods a different sort of algorithm will 'solve' the model. Numeric results will be derived by running the computational model through a metaheuristic optimization algorithm, simulated annealing, and recording the performance output and chosen design vector.

Various 'tuning parameters' for the simulated annealing algorithm and model, such as cooling rate and number of agents, will be varied. The emergence of properties such as repeatability, reliability, and optimality will be measured against real world values to validate the model. Sensitivity to changes in these tuning parameters will test the model's robustness to errors in its formulation. Moreover, qualitative dynamics of the model's response will be cataloged and used to drive conclusions regarding the optimal use of centralization in managerial command to drive performance. Results from will show that there is no single optimal approach to overseeing day to day operations in an aircraft maintenance unit. It will also show that the tuning parameters drive optimal design vector selection far more than they drive the optimal performance level itself.

CHAPTER 2: Literature Review

This chapter will review literature relevant to the optimizing managerial efforts to drive maintence unit throughput. In regards to optimization it will examine classes of optimization methods as well as the history, development, application, and advantages of simulated annealing. It will examine practices in agent based modeling and the qualities of a model that define agency. It will look into the definition, history, and benefits of decentralized organizational decision making. Finally it will provide a brief overview of factors affecting aircraft maintenance unit's performance both explicitly through case study and implicitly though analysis of contributing factors.

2.1 Metaheuristic Optimization

The first metaheuristic algorithm was likely devised by Alan Turing as a method for decrypting messages in World War II. The goal of these algorithms was to solve hard decryption problems in tractable amounts of time. They were in no way guaranteed to provide either a valid or completely correct decryption every time they were applied, but tended to provide 'quality' solutions most of the time. [3] The field of metaheuristic optimization continued from there to mature along those lines: develop methods which provide improved solutions most of the time to complex problems of the form

$$\begin{aligned}
&minimize\ f_i(\boldsymbol{x});\ \ (i = 1, 2, 3, 4 \ldots I)\\
&Subject\ to\ h_j(\boldsymbol{x}) = 0;\ (j = 1, 2, 3, 4 \ldots J)\\
&and\ g_k(\boldsymbol{x})\ \leq 0;\ (k = 1, 2, 3, 4 \ldots K)
\end{aligned}$$

for fitness function f, and constrain functions h and g, and design vector **x**. While they have diverged into different classes since the 1940s, metaheuristic methods have certain commonalities that differentiate them from non-heuristic methods. [4]

Optimization algorithms are most broadly divided into two categories: deterministic or stochastic. Deterministic methods are characterized by having deterministic outcomes, that is for any given $f_i(\mathbf{x})$ and some set of initial conditions for the design vector $\mathbf{x}$ the resulting optimal conditions and value output will be the same across any number of trials. They use non-probabilistic methods to step in the direction of local maxima or minima. Deterministic methods are farther sub-categorized into gradient based methods (such as Newton-Raphson) or gradient free methods (such as Simplex or Nelder-Mead) [4]

Metaheuristic optimization methods are more difficult to characterize. A review of some pertinent literature yields guidelines to what constitutes a metaheuristic strategy:

- Metaheuristics are high-level strategies which guide design space searches [5]
- Metaheuristics are non-deterministic and involve at least one probabilistic component, but are not purely random searches [5]
- Metaheuristics produce either locally optimal, globally optimal, or otherwise high quality solutions that simple gradient searches would not offer given a set of initial conditions [6]
- Metaheuristics are often nature inspired [4]

Their commonalities can be broken down into presenting qualities of diversification and intensification, and balancing the tradeoffs between them. Intensification, in this sense, refers to the algorithms ability to search an increasingly local area and improve upon the highest quality/most fit solution from the previous iteration. Diversification is the process by which the strategy spans the solution space to avoid being trapped in local optima. [7]

Strategies of pursing and balancing diversification and intensification then further subdivide metaheuristic methods into population-based and trajectory based method. Population based methods instantiate an initial set of solutions across the solution space. Within each iterations they produce a set of solutions each with a corresponding fitness. Methods for diversifications and intensification still are dependent on the type of algorithm employed. Genetic algorithms, for example, use a technique wherein the traits from the best solutions are combined with those of less fit solutions. This effectively guides the population in the direction of the most fit solutions. Diversification is achieved by applying random mutations-random changes to parameters within an individual's specific solution space vector- to a subset of the population. [7] Another flavor of population based algorithm is particle swarm (PSO). Intensification in PSO is achieved as each individual moves towards their most fit nearest neighbor, causing the population to cluster towards the most fit solution. Diversification is achieved by adding a degree of randomness to the path each agent takes to arrive at the new position. [8] In all cases solutions intensify around the most fit solution within the set in any given iteration.

Trajectory based methods, such as simulated annealing, use a single actor to search the solution space. New iterations may or may not deliver a solution in the neighborhood of the prior solution. Initial conditions are seeded randomly, then a process of intensification and diversification proceeds. Again, the precise method of intensification and diversification depend on the algorithm chosen. [7] In the case of simulated annealing, which will be discussed in further detail in the next section, diversification is achieved in two ways. Early in the search process the trajectory is dominated by random walk, late in the process the random perturbation and probabilistic selection allow the trajectory to escape local optima. Intensification is achieved by probabilistic rejection of worse solution and a cooling process inspired by metallurgy. [9]

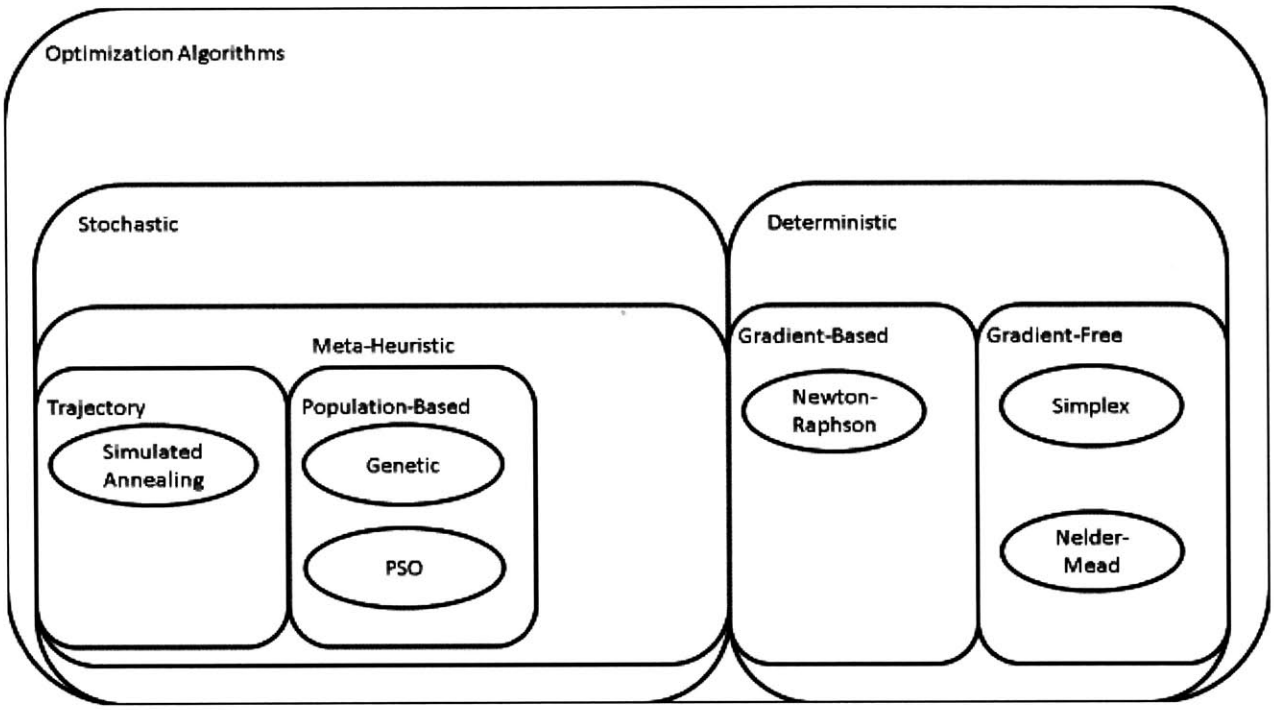

Fig 2.1: Breakout of classes of computational optimization techniques discussed or described in this section

2.2 Simulated Annealing

Simulated annealing is a class of metaheuristic optimization algorithms originally developed by Kirkpatrick, Gelatt, and Vecci in 1983. Like genetic algorithms or particle swarm optimization, simulated annealing belongs to the class of metaheuristic-sometimes called evolutionary-optimization methods. As a class metaheuristic methods are probabilistic and do not guarantee that global optimality is found, nor do they necessarily repeatedly find the same local maxima from any given initial seeding. [4] Simulated annealing is itself derived from the Metropolis algorithm for numeric simulation in many-body systems, and draws on deep connections between statistical mechanics and combinatorial or multivariate optimization. [10]

The method provides a specific set of advantages and disadvantages over other evolutionary schemes. [11]

When simulated annealing was first developed as an optimization strategy there were two dominate categories into which optimization algorithms were batched: divide and conquer or iterate and improve. Divide and conquer methods, as one might expect, subdivides problems into subproblems of manageable size, disconnects the subproblems, solves, then patches them together. If the initial problem does not decompose into naturally disjoint subproblems than the errors from patching can out weight the gains from optimization. Iterative improvement also yielded problems in complex solution spaces. As a strategy it begins in some seeded location then steps in a prearranged set of directions until an improved location is found. When no direction produces improvement, or sufficiently large improvement, the search terminates. This class of techniques typically resolves to local optima near the initially seeded 'known good configuration.' Kirkpatrick et al noted that statistical mechanics can be applied to resolve the problem of finding optimal or near optimal configurations in large, many bodied systems: ground state configurations in large bodies are very rare, but dominate annealed materials. In the case of simulated annealing, cooling processes are modeled by the Metropolis algorithm.

[10]

First published in 1953 as a method to resolve equations of state with computers using Markov Chain Monte Carlo Simulation, the Metropolis method had clear analogies to iterative improvement. The Metropolis algorithm subdivides macroscopic cooling bodies into aggregated segments of average energy. The macroscopic systems energy at any step is calculated using

$$eqn\ 2.1: E = \frac{1}{2} \sum_{\substack{i=1, \\ i \neq j}}^{N} \sum_{j=1}^{N} V(d_{ij})$$

For distance between aggregated particles d_{ij}, and potential difference between particles i and j for N particles. A particle is then perturbed within the vector space by

$$eqn\ 2.2: X_n \rightarrow X_n + \alpha\varepsilon_n$$

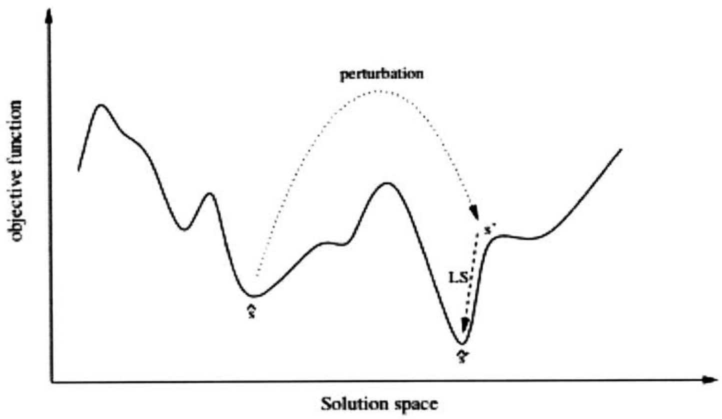

Fig 2.2: Blum & Roli depiction of perturbation [8]

For some vector **X** in n-dimensional vector space with maximum displacement α and random uniform continuous variable bounded by -1 and 1 ε. System energy is then re-evaluated by eqn 2.1. If ΔE<0, that is the systems energy has approached the ground state, the new configuration is kept and process repeated. If the perturbation yields ΔE>0 then the new configuration is kept if |ε|<exp(ΔE/kT) for system temperature T and Boltzmann constant k. Otherwise the new configuration is discarded, the prior configuration kept, and the process iterates again. [12]

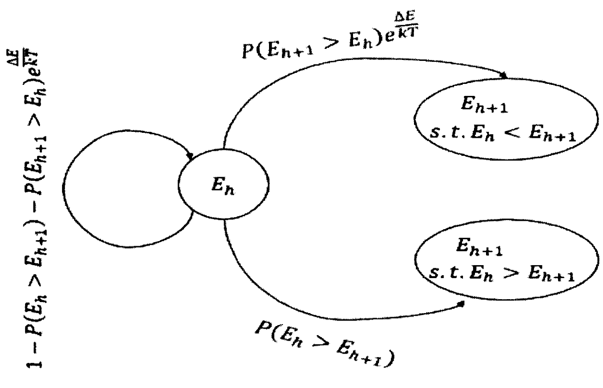

Fig 2.3: Markov chain for the hth step of the Metropolis Algorithm

Kirkpatrick et al noted that processes of rapid cooling were analogous to iterative improvement. They constantly move macroscopic bodies towards lower energy configurations but 'freeze' at metastable local energy minima. [10] The stochastic nature of the Metropolis procedure does not absolutely prevent the problem of metastable optima solutions [4], but it allows the configuration an out so long as the system's absolute temperature is positive and non-zero. In this critical way, simulated annealing as implemented through Metropolis differs from strict 'iterate and improve' strategies. This implementation allows the algorithm to cut across the iterate/divide archetype for optimization. Early on in the process the low resolutions of the features (maxima, minima, discontinuities, etc) of the eventual fate of the system are determined, as many new configurations reduce the initial aggregate energy of the system. As the system cools, a more refined, more explicitly trajectory approach takes over revealing specific details about the features. [10]

Mechanically, a simulated annealing is an adaptation of the Metropolis algorithm which can be tuned to resolve a diverse set of problems. Instead of using an explicit energy function (eqn 2.1) a situation specific cost function may be applied. Likewise, sets of parameters $\{x_n\}$ defining a decision vector replace the vector space X_n (eqn 2.2). Initial conditions are seeded randomly. Conditions for the system 'freezing' in a local or global optimal solution are control parameters within the algorithm's implementation. These control parameters can take forms such as offset from a desired system temperature needed to achieve freezing, number of reconfiguration attempts after freezing, total number of iterations, or cooling schedule. [10]

No definitive method exist to determine *a priori* what algorithm to select when exploring a new solution space. Moreover, it is not clear that for non-linear, large, global optimization problems an efficient search method exists. [13] The literature does reveal some guidelines that might lead one to select simulated annealing as a method:

1. The system under observation has few distinct, non-interchangeable parts [10]
2. The system under observation is locally amenable to hill-climbing [4]
3. The system under observation can be understood with a single objective [11]

Using these guidelines, simulated annealing appears to be an appropriate choice for the proposed research model.

2.3 Agent Based Modeling

In traditional organizational modeling, an organization is represented as an entity with a unary objective which maximizes some form of profit. However, single agent modeling with a single goal makes for an uninteresting model of an organization and betrays underlying complexity. Implementing organizational models using multiple agents provides insights that traditional systems level modeling does not. Specifically, agent based methods provide numerically driven results rather than proven results, can build complex and realistic environments that cannot be easily solved mentally, and can they can describe medium or short run dynamics rather than focus on long run equilibria. [14]

In order to implement an agent based model, the multiple agents must meet certain criteria. The model must have *agents*, *groups*, and *roles*. Agents are active communicating entities and belong to at least one group. A group is a set of agents sharing some common characteristic. While agents may or may not communicate, they may only communicate with agents who share a group. Agents also play roles, abstract functional positions. [15] Agents are endowed with objective functions which may or may not align with organizational goals and take actions with some degree of autonomy. Agents are also 'smaller' than organizations. More formally

$$J(\vec{x}) = \begin{bmatrix} T_1 \\ T_2 \\ \dots \\ T_n \end{bmatrix} \quad H_i(\vec{y}) = \begin{bmatrix} S_{i,1} \\ S_{i,2} \\ \dots \\ S_{i,k} \end{bmatrix}$$

$$\vec{y} \bigcap \vec{x} \neq \varnothing$$

$$H_i \not\subseteq J \ and \ J \not\subseteq H_i$$

Eqns 2.3: Conditions of agency

For organizational objective vector J with n objective functions T_n, agent objective vector for the i-th agent H_i with k objective functions $S_{i,k}$, and resource vectors $\vec{x}$ and $\vec{y}$. Organizations can also adapt to agents through the use of super-agent processes. [14] In particular, adaptive organizations can be modeled well with simulated annealing. [16]

2.4 Centralization & Decentralization in Decision Making

Within organizations decision making exists on a spectrum from completely centralized to completely decentralized. Internal decision making culture of firms typically lies somewhere between these two extremes, but the extremes provide specific advantages to organizations. Tuning between centralization and decentralization is dependent on a number of factors from realms of culture, capability, and cost. Tuning can also affect motivation, creativity, [17] productivity and job satisfaction [18].

The field of study regarding centralized and decentralized decision making likely began, or at least was made mainstream, with McGregor's The *Human Side of Enterprise.* In it McGregor outlines a role for management that is intrinsically different than traditional command and control. Where command and control focuses on management dictating requirements to subordinates, McGregor's new model focuses on unleashing subordinate potential. [19] He coined the two competing focuses as Theory X and Theory Y respectively. Managers who prefer one approach over the other also have strong models about their subordinates. Theory X practitioners believe that their subordinates are inherently lazy, lack self-motivation, and cannot facilitates organizational problem solving. Mangers are needed to prevent shirking and give direction to employees. Theory Y practitioners believe quite the opposite, that subordinates are self-motivated, capable, and organizationally minded. Managers, then, exist to provide the tools and contexts needed for subordinates to make decisions and solve problems. Management practices which exist in Theory Y environment like delegation, participative leadership, etc [20] are hallmarks of decentralized management.

The conditions under which an organization will prefer centralized or decentralized decisions and control are varied: personal motivation, government regulations, national cultures, individual personalities, etc. Conditions for centralization/decentralization which are internal to the organization are cost, trust, and motivation. [21] Trust and motivation play the simplest role in the dynamic. If managers do not trust subordinates then the organization will be highly centralized, while if the opposite is true decentralized networks will dominate decisions. Similarly when local decisions makers, like operators, are either unmotivated or work better when told what to do the organization will prefer to centralize. [17] More complex is the effect of communication costs on centralization. When costs of communication are exceedingly high, to the point where the cost of communicating outweighs the expected benefit of communicating, decision making is decentralized. As the cost of communication falls to the point that establishing a few lines of communication have expected net benefits organizations begin to centralize decision making. With sufficiently low costs of communication, organizations can begin to decentralize again, building a complete network map without leaning wholly on one node. [22] [17]

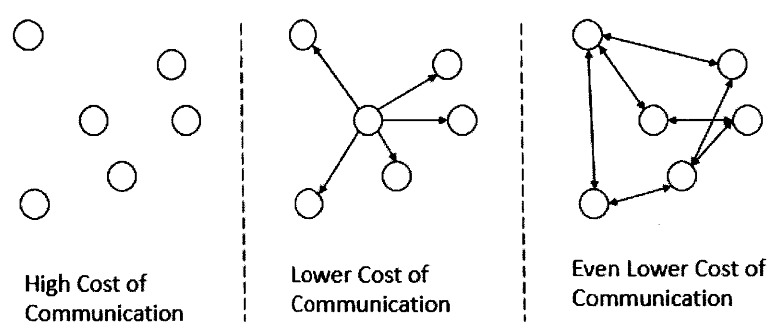

Fig 2.4: Depiction of networks based on communication cost

Centralized decision making models have merit. Centralized decision makers tend to have at least as much if not more information than decentralized counterparts and more information can lead to better decisions. Decentralization, however, still holds a number of advantages over centralized models. Participatory decision making correlates positively with both productivity and satisfaction. [18] Autonomy in general makes work more enjoyable. The ability to make decisions locally leads to improved creativity. Finally increased autonomy also improves employee motivation and in turn higher quality output. [17]

2.5 Factors Affecting Aircraft Maintenance

In the period from Jan 2000 to Dec 2011 the Korea Air Force lost 10 F-5s to crashes. This in spite of the Korean aviation system receiving the highest score in the Universal Safety Oversight Audit Program conducted by the International Civil Aviation Organization (ICAO) in 2008. Dissatisfaction with maintenance activities which led to the incidents can be viewed as a problem regarding the quality of maintenance service. While individual errors and scarcity of parts contribute to aggregate maintenance non-conformities, the largest contributors to unit quality performance are safety climate, individual health, and individual attitudes. [23] Safety climate is dictated largely by management through both degree of concern and management style. [24]

There are many physical factors that can affect operators', in this case aircraft maintainers, efficacy. Temperature, illumination, noise level, type of noise, scarcity of parts, and documentation all represent physical factors that impact efficacy. [23] However, total quality management (TQM) techniques focusing on hardware which improves physical factors is secondary and requires buy in throughout the organization. As a result processes seeking to improve efficacy have increasingly focused on social aspects. For a TQM implantation to be successful employees' attitudes towards quality must also improve. [25]

To improve quality in aircraft maintenance, first sources of errors and areas for improvement must be identified and a method of calculating effectiveness must be developed. Unit performance can be measured as an average of 4 ratios: the ratio of mission capable rate to man-hour per flying-hour rate, the rate at which aircraft are fixed against the total number of repairs to be accomplished by the next flight (sometimes called the fix rate), the rate at which repairs need to be repeated on the next flight (repeat rate), the rate at which maintenance needs to be repeated after 2 or 3 flights (recur rate), and the percent of scheduled flights carried out in a given time period.[1] Keeping a perspective of examining social factors rather than physical factors external to maintainers the factors impeding unit performance can be divided into three areas. Degradation to individual error rates, individual health, or individual attitudes towards their position/the organization all have adverse effects on unit performance. Individual errors, attitude and health are all functions of safety climate. [23]

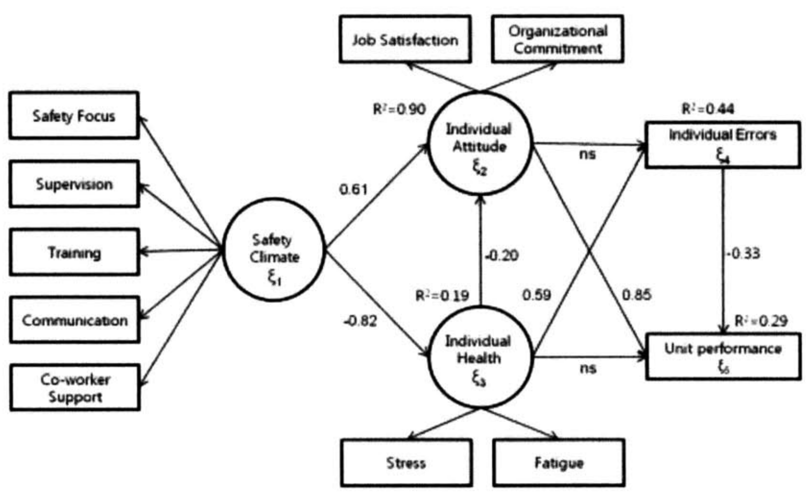

Fig 2.5: Research Model Proposed by Park et al (note, ns is not significant)

[1] These factors were cataloged by Stetz & Commenator as measures for efficiency and efficacy in an aircraft maintenance unit [39] [40]

Influencing safety climate within a maintenance community is not as simple as increasing supervisory commitment to safety though. Macroscopic safety climate is dependent upon the leadership model employed by supervisors and managers. [24] Bass has categorized and characterized leadership environments thusly:

- Constructive Leadership: a hierarchical (having reward for effort exchanges) style with an intermediate level of concern for employees welfare, also referred to as contingent reward (CR)
- Corrective Leadership: an error detection model which involves monitoring subordinates' outputs in relation to assigned standards, also referred to as management by exception (MEA)
- Laissez-faire: the lowest level of concern for employee welfare which disowns supervisory responsibility (LF)
- Transformational: a model built on value based individual interactions (TL) [26]

Zohar notes that managerial commitment to safety under each leadership model contributes differently to the prioritization of safety by subordinates as a whole. Moreover, the polarity of the change in prioritization is not necessarily the same as the polarity of the change in commitment from management. [24]

CHAPTER 3: METHODS

This chapter will present the methods used in generating and analyzing data. It will provide an exploration of the proposed multi-agent research model based on the model proposed by Park et al. It will examine the modules, calculations, and constraints of the computational model proposed. It will give a general overview of the implementation of simulated annealing used focusing on the elements that are unique to the algorithm rather than generic to all single objective simulated annealing processes: the specific perturbation algorithm, fitness function, and control parameters. It will also give an introduction to the methods of analysis to be undertaken in the following chapter.

3.1 Proposed Research Model

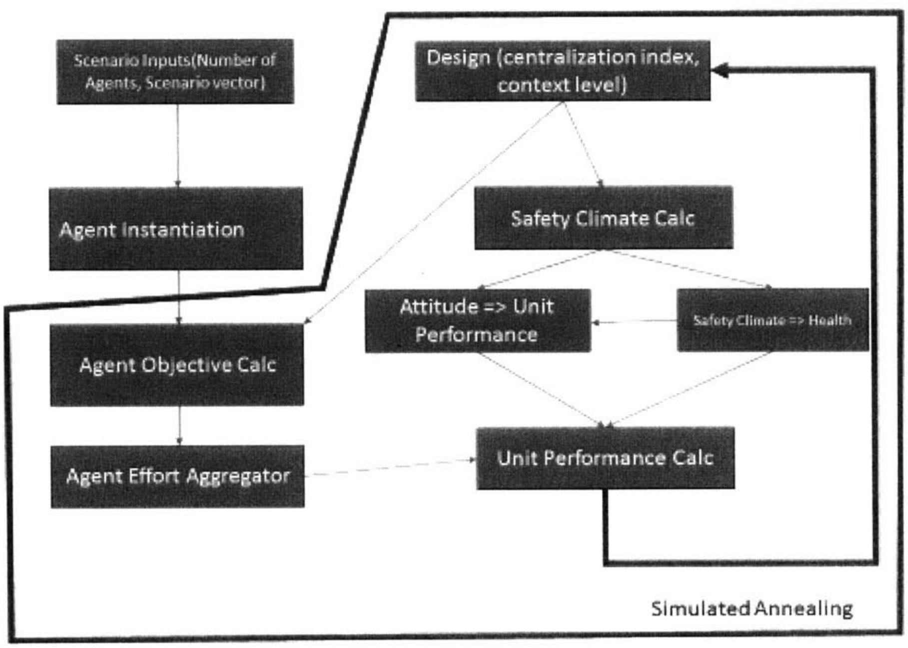

Fig 3.1 Computational model with modules

The purpose of the above model is to capture the critical effects of safety climate, individual health, individual attitudes and individual errors on unit performance as described by Park, Kang & Son. However, increasingly imposing safety climate requirements on operators can carry deleterious effects as well. The model seeks to capture the way individuals react to centralization of safety climate policy, how such activity affects their individual attitudes, and how changing group attitudes further influence individual attitudes. The following segments will discuss each modules' foundation in literature, inputs, outputs, constraints, and internal functions.

Scenarios, design inputs, and agent instantiation

The model built tests two different "scenarios" of agent compositions: a constant tolerance across all agents, and a Gaussian distribution bounded by 0 and 1. The case chosen is defined by a scenario vector that also defines the mean value of the distribution and some measure of spread. In the case of the constant tolerance the spread value is null, in the case of the Gaussian distribution the spread value is a standard deviation. Agents are then individually assign a tolerance level for centralized decision making indexed from 0 to 1.

The design values for the model are centralization and leadership model. Assigned safety priority serves as a proxy for leaderships' commitment to commanding or centralizing safety climate decisions. This is based off Zohar's measurement of assigned safety priority which asks the following questions of subordinates about their direct supervisors:

1. Do direct supervisors expect employees to 'cut corners' and work faster when behind schedule
2. Do direct supervisors ignore safety infractions if there is high schedule pressure
3. Do direct supervisors ignore safety infractions if there is no resulting injury
4. Do direct supervisors only reward production

5. Do direct supervisors get angry when they see operators performing unsafe acts [24]

Responses in agreement with question 5 and disagreement with questions 1-4 increase leadership commitment to safety and move safety trade-off decisions away from operators. Centralization is a continuous variable taking values from 0 to 1. Leadership model is a leveled variable which takes on one of Bass's four leadership modes: transformational, constructive, corrective, or laissez-faire. [26]

Safety Climate Calculation

Park et al's model defines safety climate as a high order construct which decomposed into a number of factors: safety focus, supervision, training, communication, and co-worker support. However, the model yields no method for calculating safety climate for some quantified set of these inputs. [23] Zohar proposes a method for calculating safety climate from centralization and

Leadership context and centralization are passed to the safety climate calculator which calculates safety climate as follows:

$$Eqn\ 3.1{:}\ \ Safety\ Climate = Compression * (\beta(Leadership\ Context) * Centralization + Intercept(Leadership\ Context))$$

Mappings of beta and intercept to leadership style can be found in figure 3.2. Safety climate here is measured as preventative action climate. Preventative action climate is characterized by worker responses to a 10 question survey. A subsection of these questions pertain to how supervisors interact proactively with subordinates to improve or maintain safety: do supervisors approach workers during working hours to discuss safety concerns, do supervisors take any suggestion which might improve safety seriously. [27] Compression factor here is 0.2 to correct

Zohar's 0 to 5 safety climate scale to the model's 0 to 1 scale. An intercept factor is also added to translate Zohar's 'leadership commitment' factor from 1 to 2 to our indexed 0 to 1. [24]

Leadership Style	**Beta**	**Intercept**
Transformational	0.77	2.5
Contingent Reward	0.66	3
Management by Exception (corrective)	-0.23	3
Laissez-faire	-0.34	3.5

Figure 3.2: Mapping of Leadership context to Beta and Intercept for eqn 3.1

Agent Effort Calculation and Aggregation

One of the advantages of agent based modeling over non-agent modeling is agent based models allow exploration of assumptions and integration of disparate sources. [14] There has been very little work done in the realm of quantifying how employees react to centralization of decision making. However, in so far as autonomy makes work enjoyable [17], centralization of decision making is generally characterized in economic terms as a 'bad' rather than a good for workers. [28] Workers across employment and skill spectrum require a reservation wage to begin working at all, as labor in absences of wages is also a bad. [29] Utility theory suggests that workers at the reservation wage balance the aggregate bad of the labor with the utility of wages at least to the point where they are in equilibrium. [28] We can use the notion of reservation wage and utility to model each individual agents' tradeoff between centralization and effort.

The primary tradeoffs discussed in regard to centralization are motivation and creativity. Individual agents are instantiated with a tolerance for some degree of centralization. Once that threshold has been breached they respond by reducing their output such that

$$eqn\ 3.2{:}\ 1 - Max[0, \gamma * (Centralization - Threshold)] = Individual\ Effort\ Facor\ (IEF)$$

For some tolerance to centralization γ. A γ value less than 1 represents a less dramatic response to centralization (with a value of 0 representing no response), and a γ value greater than 1 represents a more dramatic response. Total Effort Factor is then averaged from Individual Effort Factor.

$$eqn\ 3.3\text{:}\ Total\ Effort\ Factor\ (TEF) = \sum_{i=1}^{n} \frac{IEF_i}{n}$$

Individual Health, Attitude & Error Calculation

Individual health & attitude are both calculated from safety climate. Park et al define individual health as the negative mental health of an individual unit. It therefore correlates negatively with unit performance and individual attitude. The individual health score is calculated by

$$eqn\ 3.4\text{:}\ Individual\ Health = Safety\ Climate * 0.82$$

Individual health score is passed to the individual error calculator. [23]

Individual attitude is defined as the attitude an operator holds about aspects of their work environment, specifically job satisfaction and organizational commitment. It correlates positively with unit performance. Individual attitude score is calculated by

$$eqn\ 3.5\text{:}\ Individual\ Attitude = Safety\ Climate * 0.61$$

Individual attitude score is then passed to Unit Performance. [23]

Individual errors are defined as unintentional deviation from operating procedures, recommended best practices, institutional rules, and standards related to maintenance. It correlates negatively with unit performance and is calculated by

$$eqn\ 3.6: Individual\ Error = Individual\ Health * 0.59$$

Individual error score is then passed to the unit performance calculator. [23]

Unit Performance Calculation

The 'unit performance' module calculates the unit's performance based on the calculation proposed by Park et al and modified by Total Effort Factor:

$$eqn\ 3.7\ Unit\ Perf = TEF * [\alpha * individual\ attitude + \beta * individual\ health + \gamma * individual\ error$$

Their study found α to be 0.85, β to be not significant, and γ to be -0.33. [23] The individual attitude portion of the calculation is accomplished in the 'attitude to unit performance' module and passed to the 'unit performance' module while individual errors are passed directly to the module from the 'safety climate to individual error' module. The unit performance calculation is then passed into the simulated annealing function to assess fitness.

3.2 Simulated Annealing Overview

Fitness & Perturbation Function

After the unit performance is calculated for a given set of inputs (in this case centralization and leadership context) it is passed to the simulated annealing function to assess its fitness and the design variable is perturbed. Since simulated annealing functions are minimum seeking, [10] unit effectiveness score is multiplied by -1 before assessment. After the fitness is assessed and cataloged, the algorithm selects one of the two design variables at random and then perturbs it using equation 2.2.2. The process of perturbation and fitness assessment is repeated until optimality is found.

Parameters

The particular implementation of simulated annealing that the model uses has a number of control parameters: cooling schedule, cooling rate, and freeze iterations. Cooling schedule chooses between either an exponential or linear cooling schedule. Cooling rate (dT) changes the speed at which the system cools according to

$$eqn\ 3.8\text{: } Exponential\ Schedule\text{: } T_j = dT^{j*T_0} \qquad eqn\ 3.9\text{: } Linear\ Schedule\text{: } T_{j+1} = T_j - dT$$

Finally, the freeze iteration control parameter defines the number of iterations the algorithm will run at a given temperature without finding a better solution. If no improved solution is discovered within the prescribed number of iterations, the algorithm will stop and declare optimality found.

3.3 Methods of Analysis

To determine the model's merit in the pre-validation phase, post-optimality analysis will be accomplished. The principal response mechanism which will be used to conduct the analysis will be optimality results. Different instances/tunings of the model will be compared across optimal unit performance, and the design inputs. To test reliability of the simulated annealing algorithm as a method for finding optimality in this model, optimality results across independent runs of agents with constant tolerance will be compared for consistency. A similar test can be conducted for a seeded normal distribution. To overcome the stochastic nature of instantiating agents from probability distributions repeatability will also be subjected to some scrutiny. The number of agents needed to produce a reliable distribution that provides repeatable results in independent runs with constant seeding parameters (i.e. distribution type, number of agents, standard deviation) will be tested.

To compensate for the uncertainty present in any model prior to validation, the sensitivity of optimality to input parameters and model constants will be tested. In testing sensitivity, change in unit performance (optimal results) as well as change in centralization and leadership context (optimal inputs) will be measured. The model's robustness to change in these factors will be assessed. [30] A table containing all sensitive variables can be found at the end of this chapter.

Since models exist to help inform decision makers [30], and due to exigent factors those decision makers may have their hands tied into a specific leadership context, optimal results will be specifically assessed across the four leadership contexts. When conducting sensitivity analysis of any particular variable global optimal sensitivity will be assessed. Additionally optimal results constrained to only one leadership context will be subjected to sensitivity analysis. Sensitivity within each leadership context will provide insight to their robustness.

Variable Type	**Variable**	**Possible Values**	**Description**	**Symbol**
Design	Leadership Context	[TL, CR, MEA, LF]	Leadership context as described by Bass	x_0
Input Parameter	Agent Population	Integers, [1, inf)	Number of agents in the model	n
Input Parameter	Mean Tolerance	Continuous, [0 , 1]	Average tolerance for centralization	μ
Input Parameter	Distribution	[Constant Tolerance, Bounded Normal]	Distribution of tolerance to centralization	P
Input Parameter	Spread	[Constant Tolerance, Bounded Normal low std dev to high std dev]	Density of tolerance generating PDF	σ
SA Control	Cooling Schedule	[Linear, Exponential]	Method of cooling used in SA	B

SA Control	Cooling rate	Continuous, (0, 1)	Rate at which SA cools under each schedule	dT
SA Control	Freeze Iteration	Integer[1,inf)	Number of iterations without improvement required to stop at a given temperature	N
Model Parameter	Gamma	[0,inf)	Response of individual employees to centralization	γ

Fig 3.3 Summary of sensitive variables

CHAPTER 4: DATA ANALYSIS

The model under investigation has a number of stochastic elements. The optimization algorithm has a probabilistic rejection component, random initial seeding, and random perturbation. The model is itself instantiated with a random distribution of agent tolerances, guaranteeing that each individual run, even if all input parameters are the same, shows only a subset of all possible responses. The goal of this section will be to determine the model's reliability, repeatability, and robustness to various tuning parameters across both optimal result and design vector space. It will also conduct a qualitative validation against the literature.

4.1 Reliability & Repeatability Results

Reliability data were generated by running a set of input parameters through the model and cataloging the best configuration of design variables found as well as the resulting performance metric. The process was repeated 100 times for each set of tested input parameters. In these test simulated annealing control parameters where held constant at dT=0.9, N=5, with an exponential cooling schedule. Similarly, model parameters where held constant with γ=1.

The constant agent tolerances produced a mean performance of 0.485 across all runs with a standard deviation of 0.003. Minimum performance was 0.474 with design vector [.004, LF] and maximum performance was 0.487 with design vector [0.900, CR]. The model returned a contingent reward (CR) optimal leadership context in 98% of instances and a laissez-faire (LF) context in 2%. There were no instances of either transformational (TL) or management by exception (MEA) returned. Since the model returned two possible optimal leadership contexts and centralization contributes differently to safety climate, and therefore performance, measures of mean centralization across all runs have little meaning. Across runs which returned a CR context, centralization averaged 0.880 with a standard deviation of 0.03. Across runs which returned an LF context centralization average 0.005 with a standard deviation of 0.003.

Testing reliability for a normal distribution of agent tolerances yielded a mean performance across all runs of 0.474 with a standard deviation of 0.003. Difference between the maximum and minimum performance was not significant. In all runs the optimal leadership context returned was LF, issuing a recommended centralization of 0.006 with a standard deviation of 0.005.

Input Parameter Values	**Context Returned**	**Context Rate**	**Mean Performance**	**Standard Deviation Performance**	**Mean Centralization**	**Standard Deviation Centralization**
P=Constant Tolerance, σ=0, n=10, μ=.90	-	-	0.485	0.003	-	-
	CR	98%	0.485	0.003	0.880	0.03
	LF	2%	0.474	0.0001	0.004	0.003
P=Normal Bounded*, σ=0.2, n=30, μ=0.40	-	-	-	-	-	-
	LF	100%	0.474	.003	.006	.005

Fig 4.1: Summary of reliability results

Repeatability data was generated by instantiating agents with a standard set of input parameters (P=Normal Bounded, σ=.2, μ=.5) and some number of agents. The model then runs, is optimized, the results cataloged, then a new set of agents is instantiated with the same set of input parameters. The process is repeated 100 times. Repeatability in context emerges within a

2 agent model for the given inputs, recommending an LF context in all instances. Performance repeatability, as measured by the signal to noise ratio, emerged at 3 agents yielding a ratio of .0025. Due to the low centralization values presented in optimal LF context repeatability in centralization never emerges to the same degree as other parameters, but standard deviation does fall reliably below the mean centralization value by 30 agents.

In each instance the simulated annealing algorithm was able to reliably find an optimal performance. Moreover, analysis of run history data reveals that the search portion of the heuristic algorithm is spanning all hyper-planes. As an example a randomly selected history of 1 run conducting 125 searches searched the TL hyperplane 11 times, the LF (optimal in this case) hyperplane 67 times, CR 23 times and MEA hyperplane 24 times. For the constant tolerance and Normal distribution 95% of results were equal to the mean ±1.5%. The model produces repeatable results with 30 agents, a number less than the minimum size of an aircraft maintenance unit in the US Air Force. For computational ease n will be set to 35 (a squadron's minimum size) unless otherwise noted from here on out. [31]

4.2 SA Parameter Sensitivity Analysis

The simulated annealing algorithm has three controllable parameters: cooling schedule, cooling rate, and freeze iterations. Simulations were performed using 35 agents instantiated with P=Normal Bounded, σ=.2, μ=.8, and γ=1. Simulated annealing control parameters were varied across dT space from .1 to 1 for linear cooling and .1 to .9 for exponential cooling (a dT value of 1 for exponential cooling does not cool the system). Signal to noise in average performance was 1100 and 785 for linear and exponential cooling respectively. The highest ratio of signal to noise ratio was found with B=Linear and dT=0.3. Freeze iterations were then tested across a range of 1 to

10 using the control parameters found to maximize signal to noise ratio above. Signal to noise was maximized with N=6.

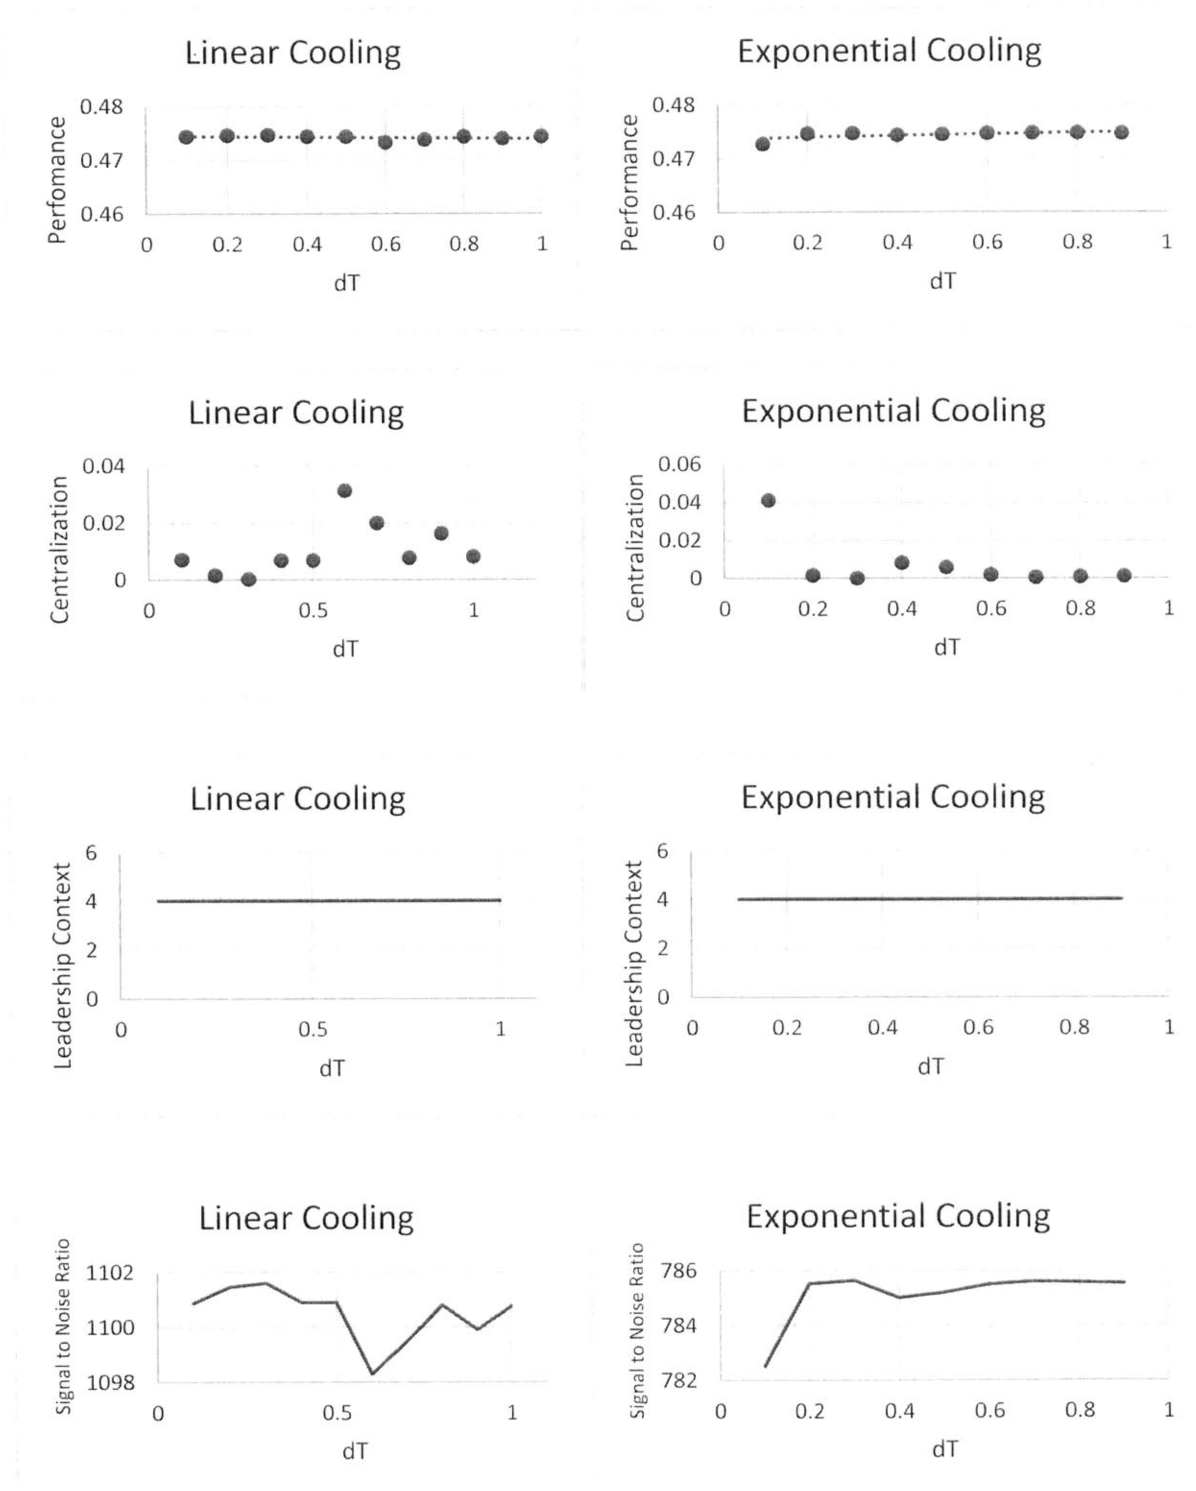

Fig 4.2: Comparison of linear and exponential cooling response across different cooling rates; leadership context mapping: 1=TL, 2=CR, 3=MEA, 4=LF

4.3 Sensitivity, Thresholds, and Emerging Effects

Having calibrated the tools for analysis, the model can now be subjected to more rigorous response analysis. As seen in the reliability & repeatability testing, each set of inputs instantiates a new sub-model with a different set of optimal decision variables. Response is not measured particularly well in terms of performance; performance is not sensitive to either changes in μ and γ. The model was run using a constant tolerance with μ and γ values varied. When average performance was plotted against γ values, changes in optimal performance were minimal with a $\frac{\delta Perf}{\delta \gamma}$ of -0.0038. Similarly, average performance was plotted against μ yielded $\frac{\delta Perf}{\delta \mu}$ of 0.0026. Across the full data set standard deviation from the mean (regardless of μ and γ inputs) was 0.0078.

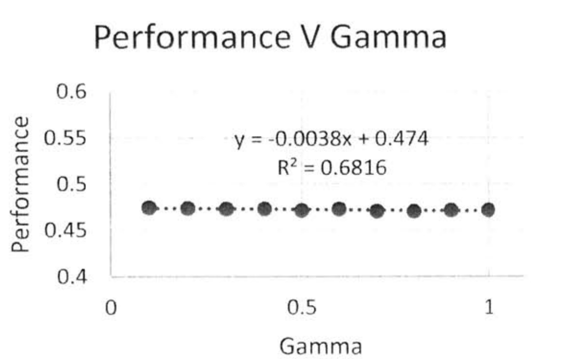

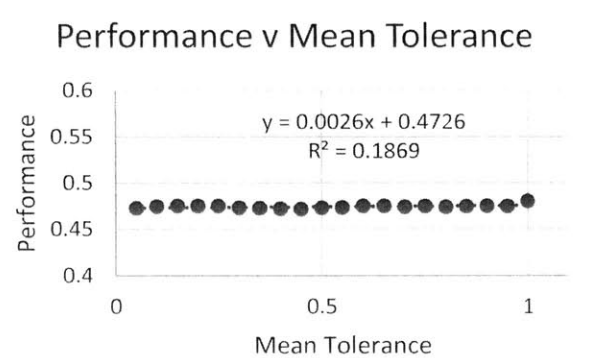

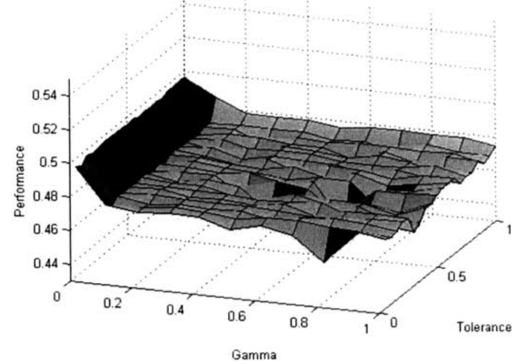

Fig 4.3: Performance against γ and μ

A more interesting view of the model's response to variation in μ and γ is to look at changes in design vectors at optimality. When gamma is equal to 0 maximum performance is simply a function of maximizing safety climate. This is achieved analytically, and verified in simulation, by choosing CR and maximizing centralization. Mean tolerance can be disregarded as the negative effects of breaching that tolerance are negated by the 0 γ value. Values of γ above zero cause leadership context chosen to be dependent on mean tolerance.

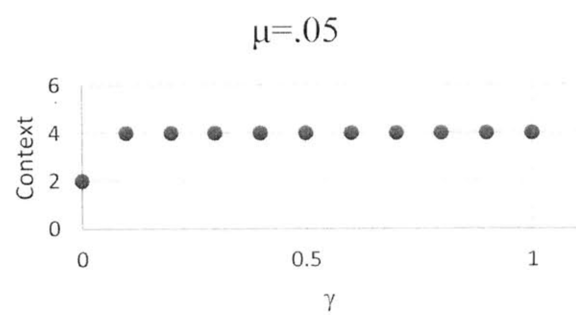

Fig 4.4

At a value of μ=0.05 the threshold for switching from CR to LF is somewhere between γ=0 and 0.1. As γ values increase, μ thresholds to switch form CR to LF increase as well. For μ values from .05 to .9 response in Context cross γ space for tested values mirrors Fig 4.4. Above that, response looks like Fig 4.5.

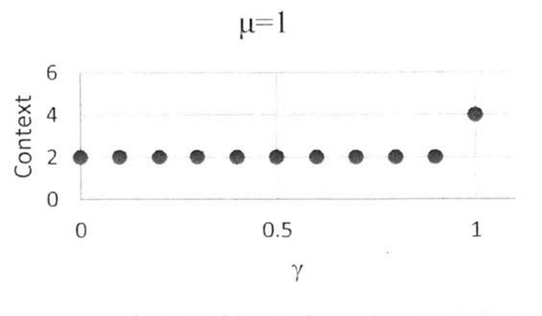

Fig 4.5

In terms of centralization submodels with lower μ value respond the similarly, they have two distinct phases. They seek to maximize centralization while CR contexts are selected (when γ values are also low), and minimize it under LF(when γ values are higher). Higher levels of μ have four phases. Under LF context, which occur at high γ values, optimality is achieved by minimizing centralization. At very near 0 values of γ and under CR context centralization is maximized. There is also a region where centralization values take on the μ value. There is an in between area, though, where employee aversion to centralization is not so large that no gains can be made from it. At the same time employee aversion to command does not permit unlimited centralization. Fig 4.6 depicts this area in black for a mean tolerance of .95 and constant tolerance. The area to the left is the region of centralization maximization, and right where centralization is set to equal the tolerance.

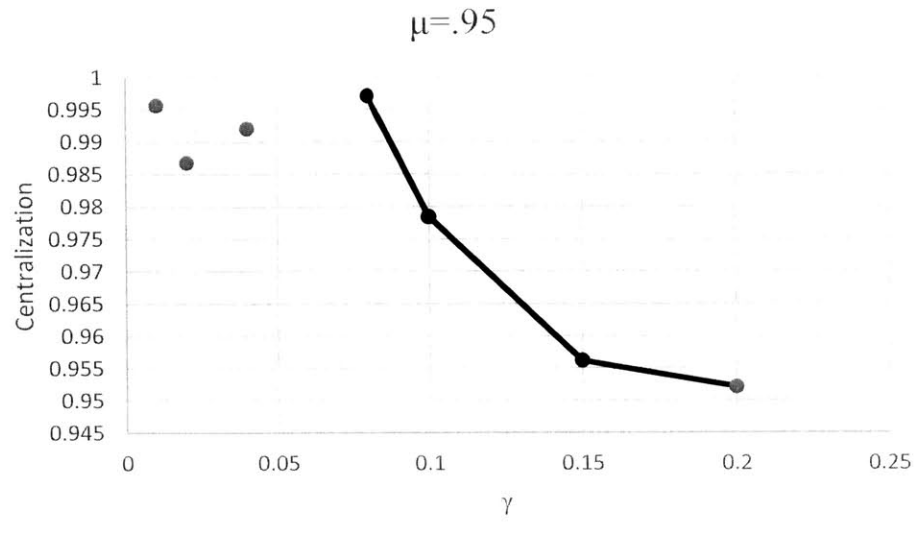

Fig 4.6

Engaging the model through agency also impacts response. Using a γ value of 0.3 and μ of .9 the model was run across a variety of standard deviation values. Agents were instantiated with a bounded normal distribution. Normally distributed agents cause the centralization and context response to shift from CR to LF event at high μ values. To ensure this effect was an not an artifact of bounding the normal distribution and thereby shifting its mean value downwards from μ, a set of constant tolerance models with μ values equal to those of each normal seeded run's resulting mean value of agent tolerance rather than the input μ value (see A-3) was run and plotted against them.

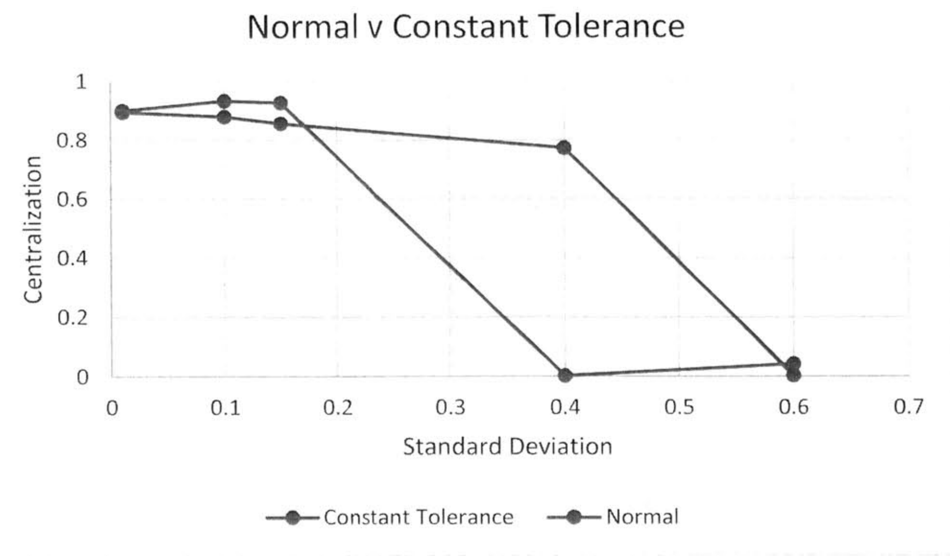

Fix 4.7: Note, constant tolerance series standard deviation values are mapped to the mean of each bounded normal distribution at that deviation, mapping can be found in A-3

Not only did agency shift to an LF recommendation earlier in terms of parameter inputs, it also shifted earlier in terms of population mean tolerance. Moreover the shift occurred in spite of median tolerance values across the population remaining reasonably flat.

4.4 Constrained Global Optimality

Natural results of the model pertaining to global optimality have explored response in swaths of the LF and CR hyperplanes. Thus far optimality conditions for the MEA and TL hyperplanes has been unexplored, as well as portions of LF and CR where they provide sub optimal responses. In tests about 87% of optimal responses returned LF and 13% returned CR as optimal. Here the other context domains will be examined for optimality, using centralization as the sole decision variable.

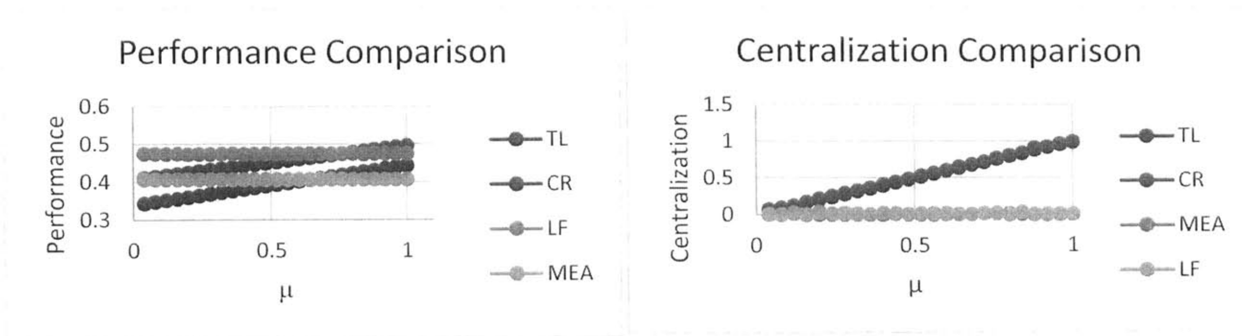

Fix 4.8: Constant tolerance varying μ with γ=0.3

Model runs whose leadership context is constrained show several interesting features. First they show the reason for LF and CR context's prevalence across optimal solutions in the unconstrained model. Across the whole spectrum of centralization any MEA or TL solution is dominated by a CR or LF one at an equal level of centralization. MEA strategies, in fact, were everywhere dominated by any LF or CR strategy regardless of underlying μ value. Repercussions of data is made more salient when considering the degree of overlap in optimal centralization between TL and CR or MEA and LF respectively. Choosing a similar value of μ and changing context from TL to CR or MEA to LF always benefits the organization in the model.

CHAPTER 5: CONCLUSIONS

"You get the steal sign, and you think, "OK, this is probably my only chance to steal in the whole game. I'm going no matter what. I'm just going to go." And then it's a slide step, you're going, and you're supposed to hold up, but you know you might not get another chance, so you keep going. It's a lot easier when you have the freedom to go on any pitch." [32]

-Coco Crisp on base stealing

5.1 Summary of Results

Literature on metaheuristic optimization and organizational modeling suggest that this form of problem is well suited to be solved by simulated annealing. The problem spans a small search area in each hyperplane, the problem has multiple hyperplanes adding an element of combinatorial difficulty, within each hyper plane and across hyperplanes the problem is 'climbable', and attempts to model an adaptive organization. [4] [10] [11] [16] Assertions from the literature are confirmed by the reliability and repeatability results. Runs seeded with the same values and deterministic agent instantiation enjoyed a high degree of convergence in performance output and design vector selection. Those runs where agents were seeded randomly (Normal) saw similar degrees of convergence across those performance and design vector selection. Randomly seeded runs also proved repeatable at realistic populations for a maintenance unit. Maintenance squadron populations have a lower bound of 35 [31] while repeatability emerged at an agent population of 30 as measured by signal to noise within performance results. Variation in the simulated annealing's tuning parameters had little effect on performance or design vector selection as well. While a specific set of parameters were chosen to minimize signal to noise ratio, the difference in that ratio between high and low values was not choosing between a noisy response and a quieter response but rather between a very quiet response (S:N of 765) and an extremely quiet response (S:N of 1100).

The model proved to be robust, in terms of optimal performance, across μ and γ values. Optimal performance at a given μ changed little as γ was varied. More interestingly, the model returned remarkably constant results in performance across varied μ. This is likely an artifact of the relative weightings for improvement in safety climate across leadership contexts with respect to centralization. Runs always attempt to maximize safety climate while not going (too far) over the boundary of tolerance for centralization. Relative weighting of gains from centralization and costs of mismanaging subordinates defines this tradeoff.

Agent based modeling, in this implementation of an operational organization, did not provide many insights. Agent instantiated runs, those that engendered each agent with a randomized objective function, were able to optimize to the same level of performance as non-agent/constant tolerance counter parts. While searching for optimality, and returning either CR or LF conditions both randomly instantiated and deterministically instantiated agents returned equal optimal centralization parameters. They differed in the μ value at which they changed context selection from one to the other. The difference between the μ at which they switch is proportional to the σ of the instantiating distribution.

Context conditional optimality did vary across μ values, with LF and CR proving to be dominant. In terms of selecting centralization factors at any give μ value, LF and MEA mimicked each other, while CR and TL were similar. This model seems to suggest that to maximize performance MEA and TL methods should be avoided out right.

5.2 Qualitative Verification

It is tempting to become bogged down in the numeric results of the simulation. The physics driving the dynamics are well documented and published, as are the weightings of dependent paths. However a number of risks, further documented in the next section, that the weightings and specific numeric response may be in doubt. The critical results for managers trying to learn from the model is the presence of dynamics and the validity of the model.

Again, the numeric results may be dubious and weightings may need to be fine-tuned. With that said, the model does conform to a qualitative validation of how aircraft maintenance responds to managerial context and command. Park, Kang & Son's model suggest that the most critical managerial (as opposed to logistic or operational) component in maximizing performance is safety climate and notes that supervision and communication play important roles in doing that. [23] The model concurs. All runs sought high levels of safety climate. In ascents of the total effort factor (γ) the model sought the maximum safety. At the opposite end of the spectrum ($\gamma=1$) safety climate was strictly limited by the maximum centralization tolerated by employees. Late μ intersection of performance curves generated by LF and CR runs in terms of safety still allowed for high safety levels to be achieved when producing optimal results. The model also concurs with notions about centralization, that independence and autonomy benefit organizations. [17] Most runs recommend a minimal (numerically 0) level of centralization and an LF context.

Dynamics, in this case meaning the various thresholds resulting in different design vector choice, reveal 4 distinct zones or recommendations for leaders and managers. If your employees have a high tolerance for centralized decision making and offer no push back, either because they do not want to or cannot, then a contingent reward strategy with maximal centralization is best suited. If the employees have other than high tolerance and offer some pushback a laissez-faire strategy with minimal managerial involvement in day to day decision making is advised. If they have a high tolerance and offer some degree of pushback adopting a CR strategy and limiting centralization to the mean level desired by all employees (barring a large spread of tolerances) is recommended. Finally, for the right balance of high tolerance and slight push back some, but not much, breaching of that tolerance threshold under a contingent reward strategy maximizes return. Where an aircraft maintenance organization, even a military one as described in the case study, fits in this landscape is dubious.

The model's ultimate conclusion echoes Coco Crisp's remarks at the beginning of the chapter. Largely employees don't want to be maimed or do harm, but sending too many signals under the wrong context can cause operators to make suboptimal decisions. Barring other exigent, operational, or logistic limitations improved performance can be eked out through better aligning managerial context and centralization with the desires of subordinates.

CHAPTER 6: RISKS & WORK TO BE DONE

6.1 Quantitative Verification and Validation

As mentioned above the model presents a view of aircraft maintenance and operations in general that are qualitatively in line with what literature has espoused on centralization versus decentralization and safety climate to performance mapping. The next step is to compare the results to recorded real word data if it exists, or to collect that data. [30] In particular, assumptions made regarding how total effort factor works should be investigated. The assumption here is that after some threshold employees respond in a linear, detracting fashion. There is also no strict empirical basis for this particular assumption, aside from the notion that when subjected to an economic 'bad,' rational people respond.

6.2 Multiobjective Organization

The model presented focuses on a short run concerned, unary organization. One that is only concerned with optimizing unit performance today. In truth aircraft maintenance units, particularly the kind described by Park, have other goals as well. They need to train new personnel, adapt to changing work, replace personnel who quit or are fired, and generally learn as organizations. While the model can provide insight into day to day operations of an aircraft maintenance unit, it does not capture a unit's need to engage in a long run maintenance strategy. Adding a second objective which captures the degree to which the organization learns, and the tradeoffs between learning and operating. [33]

Multiobjective modeling can also elucidate tradeoffs hidden within the efficacy metric. Park, So, and Kang define efficacy as the average of four metrics: the ratio of mission capable rate to man-hour per flying-hour rate, the fix rate, the repeat rate, and the recur rate. The metric obfuscates how maintenance is performed, that is whether maintenance activity is driven by operations tempo and maintenance errors. To build modules that compute each individual metric first work must be done to map safety climate as well as other exigent factors, like operational tempo or complexity, to each component of the efficacy metric.

6.3 Risks

The safety climate index is a crucial element of the model presented. Safety climate joins Park et als' efficacy calculation to Zohar's leadership criteria. Since the index is defined by two different researchers it is not clear how they map to one another: while may be safe to say that the indices 0 and 1 values are either equal or close enough for our purposes the intermediate values though may not. There is no guarantee that the mapping between the two measures is linear, nor is it necessarily the case that the mapping have constant monotonicity. A possible fix for this, in terms of sensitivity analysis, would be to add a shape factor to the safety climate calculation, allowing it to vary either systematically or randomly in simulation. Alternatively, during a quantitative V&V the effect of mis-mapping can be investigated as a possible cause of discrepancies between the model and reality.

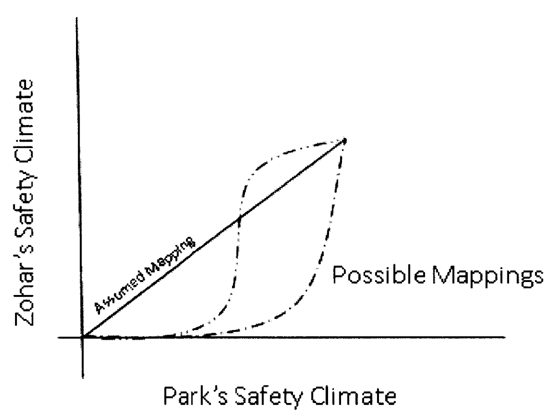

Fig 6.1: Possible mappings of safety climate index

The weightings given in the case study are based on surveys of the Korean Air Force. It is doubtful that the weightings hold much relevance outside of that organization. Cross price elasticity for similar goods can vary widely across borders. [34] [35] [36] To the extent that causal paths in Park et al model are a function of managements demand for safety, subordinates demand for lower stress, the publics' and other overseers' demand for flying hours, etc there is no reason to expect similar weighting in other military aircraft maintenance units let alone in the civilian space.

Organizations certainly use leadership styles which are neither laissez–faire nor based on contingent reward. [37] Management by exception and transformational leadership have not succumb to market arbitrage but do not present in the optimized results. Why? A multiobjective approach may sort out this apparent contradiction. Transformational leadership has long been associated with learning organizations [38]. There may be subtleties in the ways that employees react to under management, for example some kind of confusion factor, that lead managers to use forms of management that do not seem optimal here.

Bibliography

[1] "Airline Maintenance Cost Executive Summary," International Air Transportation Association, 2011.

[2] "Why Has the Cost of Fixed Wing Aircraft Risen?," The Rand Corporation, 2008.

[3] B. J. Copeland, The Essential Turing, Oxford University Press, 2004.

[4] X.-S. Yang, "Review of Metaheuristics and Generalized Evolutionary Walk Algorithm," *Int. J. Bio-Inspired Computation,* vol. 3, no. 2, pp. 77-84, 2011.

[5] T. Stützle, "Local Search Algorithms for Combinatorial Problems - Analysis, Algorithms and New," *Dissertations in Artificial Intelligence,* vol. 220, p. 203, 1999.

[6] S. Voss, I. Osman and C. Rouciarol, Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization, Norwell, Ma: Kluwer Academic Publishers, 1999.

[7] C. Blum and A. Roli, "Metaheuristics in Combinatorial Optimization: Overview and Conceptual Comparison," *ACM Computer Surveys,* vol. 35, pp. 268-308, 2003.

[8] J. Kennedy and R. Eberhart, "Particle Swarm Optimization," *Preceedings of the IEEE International Conference on Neural Networks,* pp. 1942-1948, 1995.

[9] V. Cerny, "A thermodynamical approach to the travelling salesman problem: An efficient simulation," *Journal of Optimization Theory and Applications,* vol. 45, no. 1, pp. 41-51, 1985.

[10] S. Kirkpatrick, C. D. Gelatt and M. P. Vecchi, "Optimization by Simulated Annealing," *Science,* vol. 220, no. 4598, pp. 671-680, 13 May 1983.

[11] D. Nam and C. H. Park, "Multiobjective Simulated Annealing: A Comparative Study of Evolutionary Algorithms," *International Journal of Fuzzy Systems,* vol. 2, no. 2, pp. 87-97, 2000.

[12] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller and E. Teller, "Equation of State Calculations by Fast Computing Machines," *The Journal of Chemical Physics,* vol. 21, no. 6, pp. 1087-1092, 1953.

[13] D. H. Wolpert and W. G. Macready, "No Free Lunch Theorems for Optimization," *IEEE Trans. on Evolutionary Computing,* vol. 1, no. 1, pp. 67-82, 1997.

[14] M.-H. Chang and J. E. Harrington Jr, "Agent Based Modeling of Organizations," in *Handbook of Computational Economics*, vol. 2, L. Tesfatsion and K. L. Judd, Eds., Elsevier B.V., 2006, pp. 1274-1337.

[15] J. Ferber, O. Gutknecht and F. Michel, "From Agents to Oganizations: An Organizational View of Multi-Agent Systems," in *Agent Oriented Software Engineering IV*, Springer Berlin Heidelberg, 2004, pp. 214-230.

[16] K. M. Carley and D. M. Svoboda, "Modeling organizational adaptation as a simulated annealing process," *Sociological Methods & Research,* vol. 25, pp. 138-168.

[17] T. W. Malone, "Is Empowerment Just a Fad? Control, Decision Making, and IT," *Sloan Management Review,* pp. 23-35, January 1997.

[18] K. M. Miller and P. R. Monge, "Participation, Satisfaction, and Productivity: A Meta-Analytic Review," *Academy of Management Journal,* vol. 29, no. 4, pp. 727-253, 1986.

[19] R. E. Kopelman, D. J. Prottas and A. L. Davis, "Douglas McGregor's Theory X and Y:Toward a Construct-valid Measure," *Journal of Managerial Issues,* vol. 20, no. 2, pp. 255-271, 2008.

[20] D. M. McGregor, The Human Side of Enterprise, New York, NY: McGraw-Hill, 1960.

[21] P. J. Dimaggio and W. W. Powell, "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Field," *American Sociological Review,* vol. 48, pp. 147-160, April 1983.

[22] E. V. Hipple, ""Sticky Information"" and the Locus of Problem Solving: Implications for Innovation," *Management Science,* vol. 40, no. 4, pp. 429-439, April 1994.

[23] H. Park, M. J. Kang and S. Son, "Factors Affecting Qualityand Performance- a Case Study of Korean Aircraft Maintenance Unit," *Total Quality Management,* pp. 197-219, February 2012.

[24] D. Zohar, "The effects of leadership dimensions, safety climate, and assigned priorities on minor injuries in work groups.," *Journal of Organizational Behavior,* vol. 23, no. 1, pp. 75-92, February 2002.

[25] J. C. Bou and I. Beltran, "Total Quality Management, High-commitment Human Resource Strategy and Firm Performance: An Emperical Study," *Total Quality Management,* vol. 16, no. 1, pp. 71-86, January 2005.

[26] B. M. Bass, Bass & Stogdill's Handbook of Leadership, New York: Free Press, 1990.

[27] D. Zohar, "A Group-Level Model of Safety Climate: Testing the Effect of Group Climate on Microaccidents in Manufacturing Jobs," *Journal of Applied Psychology,* vol. 85, no. 4, pp. 587-596, 2000.

[28] D. D. Friedman, Price Theory: An Intermediate Text, South-Western Publishing Co, 1990.

[29] S. L. Rynes, B. Gerhart and K. A. Minette, "The Importance of Pay in Empoyee Motivation: Discrepancies Between What People Say and What they Do," *Human Resource Management,* vol. Winter, pp. 381-394, 2004.

[30] K. W. Hipel, D. M. Kilgour and S. Rajabi, "Operations Research and Refinement of Courses of Action," in *Handbook of Systems Engineering & Management*, A. P. Sage and W. B. Rouse, Eds., John Wiley & Sons, 1999, pp. 1078-1118.

[31] *AFI38-101-Manpower and Organization,* 2011.

[32] J. Keri, "Grand Theft Baseball," 20 March 2013. [Online]. Available: http://grantland.com/features/jonah-keri-art-stealing-bases/. [Accessed 5 Dec 2014].

[33] N. P. Repenning, "Understanding Fire Fighting in New Product Developement," *Journal of Product Innovation Management,* vol. 18, no. 5, pp. 285-300, 2001.

[34] J. C. Cooper, "Price Elasticity of Demand for Crued Oil: Estimates for 23 Countries," *OPEC Review,* March 2003.

[35] "Cross-Price Elasticities of Demand Across 114 Countries," USDA, 2010.

[36] P. Fajnzylber and W. F. Maloney, "How Comparable are Labor Demand Elasticities across Countries?," World Bank, 2001.

[37] K. B. Lowe, K. G. Kroeck and N. Sivasubramaniam, "Effectiveness Correlates of Transformational and Transactional Leadership: A Meta-Analytic Review of the MLQ Literature," *The Leadership Quarterly,* vol. 7, no. 3, pp. 385-415, 1996.

[38] A. Coad and A. J. Berry, "Transformational leadership and learning orientation," *Leadership & Organization Development Journal,* vol. 19, no. 3, pp. 164-172, 1998.

[39] L. J. Stezt, "An ex post facto analysis of E-3 maintenance indicators in the 552nd air control wing since reorganization under an aircraft generation Squadron concept," Air Force Institute of Technology, Wright-Patterson AFB, 1999.

[40] M. A. Commenator, "Aircraft maintenance performance: The effects of the functional decentralization of on-equipment maintenance," Air Force Institute of Technology, Dayton, OH, 2001.

Appendix

A-1: Zohar, 2000 Questionnaire

Table 1

Rotated Principal-Components Analysis Factor Structure of the Group Safety Climate Scale With Items Rearranged by Factor (N = 534)

Item	Factor 1: action	Factor 2: expectation
1. My supervisor says a good word whenever he sees a job done according to the safety rules.	0.738	0.095
2. My supervisor seriously considers any worker's suggestions for improving safety.	0.776	0.125
3. My supervisor approaches workers during work to discuss safety issues.	0.738	0.166
4. My supervisor gets annoyed with any worker ignoring safety rules, even minor rules.	0.605	0.180
5. My supervisor watches more often when a worker has violated some safety rule.	0.620	0.255
6. As long as there is no accident, my supervisor doesn't care how the work is done (R)	0.077	0.676
7. Whenever pressure builds up, my supervisor wants us to work faster, rather than by the rules (R).	0.286	0.673
8. My supervisor pays less attention to safety problems than most other supervisors in this company (R).	0.083	0.611
9. My supervisor only keeps track of major safety problems and overlooks routine problems (R).	0.271	0.642
10. As long as work remains on schedule, my supervisor doesn't care how this has been achieved (R).	0.099	0.493
Eigenvalue	2.623	2.091

Note. Items are translated from Hebrew. (R) = reversal of item scores.

J. of Applied Psychology Vol 85, No.4 pg 591

A-2: Model Objectives and Constraints

Unit Performance Objective:

Maximize[TEF(0.85*IndivAttitude - 0.33*IndivError)]*

Unit Level Constraints:

Centralization>0

1-Centrailzation>0

Context=[LF, MEA, CR, TL]

Agent Response Function:

$IEF_i = 1\text{-}Max[0, \gamma * (Centralization - Threshold_i)]$

$TEF = \sum_{i=1}^{n} \frac{IEF_i}{n}$

Agent Constraints:

$IEF_i > 0$

$1 - IEF_i > 0$

$Threshold_i > 0$

$1\text{-}Threshold_i > 0$

A-3: Mapping for Fig 4.7

Std Dev Corresponding to Each Value	Resulting Mean Value of Normally Distributed Agent Tolerance	Resulting Median Value of Normally Distributed Agent Tolerance	Agent Tolerance of Constant Tolerance Runs
0.01	0.901	0.902	0.901
0.1	0.8857	0.8958	0.8857
0.15	0.8595	0.9039	0.8595
0.4	0.7766	0.9347	0.7766
0.6	0.7087	0.8524	0.7087