



CONTINUOUS DECISION SUPPORT

DISSERTATION

Jeremy P. Hendrix, GS-14, DAFC

AFIT-ENV-DS-15-D-018

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

**DISTRIBUTION STATEMENT A:
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED**

The views expressed in this dissertation are those of the author and do not reflect the official policy or position of the United States Air Force, the Department of Defense, or the United States Government.

This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENV-DS-15-D-018

CONTINUOUS DECISION SUPPORT

DISSERTATION

Presented to the Faculty
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy in Systems Engineering

Jeremy P. Hendrix, B.S., M.S.

GS-14, DAFC

December 2015

DISTRIBUTION STATEMENT A:
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

CONTINUOUS DECISION SUPPORT

Jeremy P. Hendrix, B.S., M.S.
GS-14, DAFC

Committee Membership:

David R. Jacques, PhD (Chairman), Jeffery D. Weir, PhD (Member), Bradley J. Ayres,
PhD (Member)

Adedeji Badiru
Dean, Graduate School of Engineering
and Management

Abstract

Organizations are often faced with portfolio construction efforts that require them to select one or more alternatives, subject to resource constraints, with the aim of achieving the maximum value possible. This is a well-defined problem with a number of analytically defensible approaches, provided the entire set of alternatives is known when the decision event takes place. Less well treated in the literature is how to approach this problem when the entire set of alternatives is unknown, as when the alternatives arrive over time. This change in the availability of data shifts the problem from one of identifying an optimal subset to one in which a series of smaller decisions are undertaken regarding the acceptability of each alternative as it presents itself.

This work expands upon a methodology known as the Triage Method. The original Triage Method provided a screening tool that could be applied to alternatives as they presented themselves to determine if they should be accepted for further study, rejected out of hand, or held pending until later date. This decision was made strictly upon the value of the alternative and with no consideration of its cost. Two extensions to the Triage Method are offered which provide a capability to consider cost and other resource requirements of the alternatives, thus allowing a move from simply screening to portfolio selection. Guidelines are presented as to when each of these extensions is best employed, a characterization of the performance tradeoff between these and more traditional methodologies is developed, and insight and techniques for setting the value of parameters required by the extensions are provided.

For Carole, Cameron, and Lochlan.

Acknowledgments

This work would not have been possible without the support and patience of Mr. Rich Moore and Dr. Ross Jackson in HQ AFMC/A9A. My A9A co-workers also provided valuable insights and often acted as sounding boards throughout the process. I would particularly like to thank Mr. Roger Moulder for his insights on design of experiments, Dr. Lance Champagne for his statistical guidance, and Dr. Brad Boehmke for his all-around help on navigating the Ph.D. process. I am indebted to Dr. Ted Lewis of AFLCMC/OZA for first introducing me to this topic via our work together with the AFMC Development Planning process and later supporting my research.

It goes without saying that no one successfully completes a Ph.D. program without the support and guidance of a quality committee. Dr. Jacques, Dr. Weir, and Dr. Ayres were supportive and patient throughout the (long, long) process of completing this program, and I would not be here without them.

Jeremy P. Hendrix

Table of Contents

	Page
Abstract	iv
Dedication	v
Acknowledgments	vi
Table of Contents	vii
List of Figures	xi
List of Tables	xiii
List of Abbreviations	xiv
List of Symbols	xv
I. Introduction	1
1.1 Background	1
1.1.1 The Prospect of Multiple Decision Epochs	2
1.1.2 Availability of Alternatives and Data	3
1.1.3 A Finite Time Horizon	5
1.2 Problem Statement	6
II. Background and Related Work	9
III. The Triage Method	17
3.1 Introduction	17
3.2 Principles of the triage method	22
3.2.1 Decisions over time	24
3.2.2 Cost considerations	24
3.3 The triage method	25
3.3.1 Hypothetical decision alternatives	26
3.3.2 Accounting for uncertainty with weight intervals	28
3.3.3 Cutoff value and comparison strategy	31
3.4 Illustrative example	32
3.4.1 Brief History of JIEDDO	33

	Page
3.4.2 Triage Method Applied to JIEDDO Model	34
3.5 Summary	39
3.6 Theorem Proof	41
IV. Continuous Decision Support	43
4.1 Introduction	43
4.1.1 The Prospect of Multiple Decision Epochs	44
4.1.2 Availability of Alternatives and Data	45
4.1.3 A Finite Time Horizon	47
4.2 Background and Related Work	47
4.3 Methodological Extensions	51
4.4 An Illustrative Example	56
4.4.1 Triage Results	61
4.4.2 Triage+ Results	62
4.4.3 Triage++ Results	64
4.5 Conclusions and Future Work	67
4.6 Theorem Proof	68
V. The Effects of Decision Problem Parameters on the Effectiveness of the Extended Triage Method	70
5.1 Introduction	70
5.2 Background	73
5.2.1 The Triage Method	73
5.2.2 The Triage Extensions	74
5.2.3 The Development Planning Model	77
5.2.3.1 Resulting Methods	79
5.2.4 Experimental Approach	81
5.3 Factor Effects	83
5.4 Effectiveness of Extensions	87
5.4.1 Impact of Departure Rate	87
5.4.2 Impact of α	89
5.4.3 Risk	92
5.5 Configuration Guidelines	94
5.5.1 Departure Rate	95
5.5.2 Model Size	95
5.5.3 Sensitivity Intervals	96
5.5.4 CDF Approximations	96
5.6 Future Work	97
5.6.1 Defining α	97
5.6.2 Robust CDF Estimates	97

	Page
5.6.3 Value of Information	98
VI. The Elicitation and Application of Decision Problem Parameters in the Extended Triage Method	99
6.1 Introduction	99
6.1.1 The Triage Method	100
6.1.2 The Extended Triage Method	101
6.1.3 Key Issues	102
6.2 Background	103
6.2.1 The Development Planning Dataset	103
6.2.2 Elicitation Methods	104
6.2.2.1 α Value Elicitation	105
6.2.2.2 CDF Elicitation	106
6.2.3 Application Methods	107
6.2.3.1 Varying α on a Schedule	108
6.2.3.2 Varying α Reactively	109
6.2.3.3 Varying the Calculation of $\mathbf{R}$	110
6.2.3.4 Traditional Methods	111
6.3 Results	111
6.3.1 Elicitation Results	111
6.3.1.1 α Elicitation	111
6.3.1.2 CDF Elicitation	113
6.3.2 Simulation Results	113
6.3.2.1 Fixed α	113
6.3.2.2 α Varied by Schedule	115
6.3.2.3 α Varied Reactively	117
6.3.2.4 New Calculation of $\mathbf{R}$	120
6.4 Conclusions	126
6.4.1 The Need for Historical Data	127
6.4.2 The Modified Calculation of $\mathbf{R}$	127
6.4.3 Varying α	128
VII. Conclusions and Suggestions for Further Research	130
7.1 A Definition of Continuous Decision Problems	130
7.2 Effective Methods for Continuous Decision Problems	130
7.3 Qualifying Criteria	131
7.4 Tradeoff Characterization	131
7.5 Guiding Principles	132
7.6 Future Research	132

	Page
Bibliography	134

List of Figures

Figure	Page
3.1 The Triage Method in a two-stage decision process at time t	20
3.2 The Triage Method Hierarchy	35
3.3 Triage Method Results on JIEDDO Proposals	39
4.1 Baseball Card Value Hierarchy	57
4.2 Value Curve for Baseball Card Book Value	58
4.3 Value Curve for Baseball Card Age	58
4.4 Value Curve for Baseball Card Condition	59
4.5 Triage Portfolio Value	62
4.6 Triage Net Profit/Loss	63
4.7 Triage+ Portfolio Value	63
4.8 Triage+ Net Profit/Loss	64
4.9 Triage++ Portfolio Value	65
4.10 Triage++ Net Profit/Loss	66
4.11 Triage+ Combined Performance	66
4.12 Triage++ Combined Performance	67
5.1 Average Score (95% Confidence Interval) versus Model Size	85
5.2 Performance of Triage+ Methods	88
5.3 Triage++ Performance	89
5.4 Triage+ Performance	90
5.5 Triage++ Performance	91
5.6 ++Resources and Come and Go	92
5.7 Dollar Commitment versus Time	93
5.8 FTE Commitment versus Time	93

Figure	Page
5.9 Comparison of CDF Approximations	95
6.1 α Values for Selected Values of ρ	109
6.2 Elicited CDFs	114
6.3 Fixed α Values	114
6.4 Increasing α by Schedule	116
6.5 Decreasing α by Schedule	117
6.6 +Resources Reactive	118
6.7 ++Resources Reactive with Pearson-Tukey CDF Approximation	120
6.8 ++Resources Reactive with Empirical CDF	121
6.9 Fixed α Values with New R Calculation	122
6.10 α Values Varied by Schedule with New R Calculation	123
6.11 +Resources Reactive with New R Calculation	125
6.12 ++Resources Reactive with New R Calculation	126

List of Tables

Table	Page
3.1 The Triage Method Notations and Descriptions	27
3.2 Sample JIEDDO Hierarchy Weights and Weight Space	36
3.3 Triage Method Results on JIEDDO Proposals as t increases from 0 to 6.	38
5.1 Distribution Approximations	82
5.2 Experimental Design	83
5.3 Factor Effects	84
5.4 Selection Commonality	86
6.1 Resources Applied to Fictional Alternatives	111
6.2 Resources Applied to Minimally Acceptable Alternative	112
6.3 Elicited CDF Points	113
6.4 Performance Summary	128

List of Abbreviations

Abbreviation		Page
AFMC	Air Force Materiel Command	5
DP	Development Planning	5
LP	linear programming	5
PMJ	professional military judgment	5
R&D	research and development	9
NPV	net present value	12
JIEDDO	Joint Improvised Explosive Device Defeat Organization	32
IED	improvised explosive device	33
BAA	Broad Area Announcement	33
JCAAMP	Joint IED Defeat Capability Approval and Acquisition Management Process	34
CDF	cumulative distribution function	53
FY	fiscal year	77
FTE	full-time equivalent	78
HQ	Headquarters	103

List of Symbols

Symbol	Definition
A	Hypothetical decision alternative
α	Cutoff value for the Triage Method
α_0	Highest cutoff value in the Triage Method: $\alpha_0 = \underline{V}(D)$
α_T	Lowest cutoff value in the Triage Method: $\alpha_T = \underline{V}(N)$
α_t	Cutoff value at time t in the Triage Method: $\alpha_T + f(t)(\alpha_0 - \alpha_T)$
α_+	Cutoff value for the Triage+ Extension
α_{++}	Cutoff value for the Triage++ Extension
C	Hypothetical resource cost of an alternative
C_{crit}	Hypothetical resource cost of an alternative for a critical resource
D	Hypothetical “desired” alternative with measures $\{d_1, d_2, \dots, d_n\}$
$\delta(A)$	Maximum value increase/decrease achievable within a given weight space
$f(T)$	Monotonically decreasing time function with $f(0) = 1$ and $f(T) = 0$
γ	The adjustment increment $\subset [0, 1]$ in the reactive Triage+ and Triage++ Methods
λ_{arr}	The average number of alternative arrivals per day
λ_{dep}	The average number of alternative departures per day
N	Hypothetical “needed” alternative with measures $\{n_1, n_2, \dots, n_n\}$
R_{max}	The maximum value of the ratio $\frac{V(A)}{C_{crit}} \cdot 1000$
R_{min}	The minimum value of the ratio $\frac{V(A)}{C_{crit}} \cdot 1000$
ρ	The exponential constant
S	The size of a value model, equivalent to the number of measures the model contains
T	Time constant to complete screening
t	Current time period: $t = 0, 1, 2, \dots, T$
$Tr(A_t, \alpha_t)$	Triage function with inputs A_t and t

Symbol	Definition
$V(A)$	The value score achieved by alternative A
$V_{min}(A)$	The minimum value score alternative A can achieve in a given weight space
$V_{max}(A)$	The maximum value score alternative A can achieve in a given weight space
W	The set of weights specified for a value model
W_{max}	The set of weights within a given weight space that maximize an alternative's value
W_{min}	The set of weights within a given weight space that minimize an alternative's value

CONTINUOUS DECISION SUPPORT

I. Introduction

1.1 Background

Howard classically defines a decision as “an irrevocable allocation of resources.” [45] This general definition is broad enough to countenance any number of complicating factors such as temporal issues, resource constraints, interactions between alternatives, etc. Other authors add more specificity to their definition of a decision without limiting the scope of these complications. Skinner, for example, extends the definition to “a conscious, irrevocable allocation of resources with the purpose of achieving a desired objective.” [79] As others refine the definition however, the scope begins to constrict. Lootsma adds more specificity, defining a decision as “a choice out of a number of alternatives made in such a way that the preferred alternative is the ‘best’ among the possible candidates.” [66] Notice that we now are presented with a finite set of alternatives from which to choose. Kirkwood follows this example, eschewing a formal definition altogether and instead detailing aspects of a decision: available alternatives, differing outcomes, and uncertainty as to which outcomes are associated with the alternatives. [59]

As these definitions move from general to more and more specific, they begin to more closely resemble the activity that most people associate with decision-making: choosing from a set of known alternatives. As broad as Howard’s definition of a decision is, Keeney and Raiffa state that decision *analysis* “is designed to help the individual make a choice among a set of prespecified alternatives.” [52] But does this represent the full spectrum of decisions that decision makers are called upon to make? Arguably no, it merely

represents the most analytically tractable of common decision situations. The situation can be complicated by a number of factors including, though certainly not limited to:

- The decision maker may seek to choose more than one alternative
- The decision maker may be held to multiple resource constraints
- The decision maker may have the alternative to defer a decision
- The decision maker may expect other alternatives will present themselves later

It is not uncommon for a decision maker to face a situation that is complicated by all of the above factors. While the decision analysis literature features techniques to address the first two factors above, [37, 58, 69] the last two, and particularly the final one, are largely unaddressed. The presence of these two complicating factors indicates what we term a *continuous* decision problem. Such problems are characterized by the following conditions:

1.1.1 The Prospect of Multiple Decision Epochs.

A key difference between a continuous decision situation and a traditional one is the decision maker's reasonable expectation that the ultimate decision will be achieved via a series of multiple, smaller decisions rather than a single, monolithic decision. In a continuous decision problem, the decision maker has the option to select from the available alternatives or to defer a decision with the clear expectation that they will revisit the problem in the future. The decision maker may choose from the set of known alternatives at time t , or defer a decision until $t + x$. At time $t + x$, the decision maker again may choose from the known alternatives, or defer until $t + y$, $y > x$. The expectation is that the decision environment will be more favorable at that time. Perhaps new and better alternatives will be available, or the uncertainty surrounding the known alternatives will have decreased. Perhaps the decision maker will have a better understanding of the requirements associated

with the decision problem and can thus make a more informed choice. Or perhaps more resources will have become available, rendering previously infeasible alternatives feasible. Of note, the passage of time may preclude previously available alternatives from being pursued. Whatever the motivation, the expectation of revisiting the problem marks the distinction.

This is a subtle difference, and it may be argued that this is no different from a series of traditional “one-time” decisions and may be approached as such with traditional methods. Most traditional methods however are geared toward a definition similar to Lootsma’s. If presented with a collection of inferior alternatives, they will help choose the least inferior, but may not guide the decision maker to defer. [38] The decision may be revisited later, but if the selected alternative is changed, the resources expended to date may well have been wasted. Frequently engaging in such revisions may prove costly.

1.1.2 Availability of Alternatives and Data.

Whether the decision maker is seeking a single feasible alternative or to assemble a collection that is in some sense optimal, the difficulty is compounded if there is a reasonable expectation that the entire set of alternatives is not available, or that substantially more information about the decision problem will become available at some point in the future. Indeed, this is in all likelihood the reason the decision maker might choose to defer a decision. While the expectation may not, in the end, be realized, its presence is another defining condition, altering the decision problem from a traditional choice to a continuous decision problem.

The introduction of this element makes the distinction between analysis of traditional decision problems and continuous ones more clear. As stated earlier, traditional methods of decision analysis tend to assume that the complete set of alternatives is defined, and that the decision maker only needs assistance in identifying the best option or set of options.

Continuous decision problems require analytical techniques that can provide insight to decision makers in the absence of this assumption.

Consider the classic case of constructing a portfolio of research projects. [62] In a traditional approach, the decision maker would collect potential projects until some deadline, and then evaluate them all via their preferred methodology. Resource constraints could be treated via linear programming techniques to select a subset that was optimal according to some measure. Once selected, the portfolio would be considered complete. The decision maker may then collect new proposals until a future date, at which time they could choose to evaluate the new proposals to construct an additional portfolio, or re-evaluate the entire set of known proposals. If the entire set is re-evaluated, the decision maker may face the prospect of choosing between abandoning a previously started project or accepting the “sub-optimality” of continuing.

In a continuous approach, the decision maker would evaluate each proposal as it was presented, making a series of smaller decisions about each individual project. At each decision epoch then, the decision maker is faced with not only evaluating the merits of the new alternative, but considering the possibility that a better alternative might become available later. There is still the possibility that previously selected alternatives may no longer be part of the optimal set, but since the decision maker is only considering a single new alternative at a time, the decision becomes a smaller (though not necessarily easy) tradeoff to consider.

While the resulting portfolios will likely be different, the ultimate goal of an optimal portfolio remains unchanged. Almost by definition, the portfolio constructed via the continuous framework will be “less optimal” than one constructed when the entire set of alternatives is known. This is to be expected, as decisions made in the traditional construct benefit from more complete information. The level of decrease acceptable to the decision

maker will depend on their risk attitude. [33] The desire on the part of the decision maker, and the aim of the analyst, is to gain sufficient flexibility to offset the loss of optimality.

1.1.3 A Finite Time Horizon.

While theoretically a decision may be deferred indefinitely, for practical applications an alternative or alternatives must eventually be chosen. Further, the existence of a finite time horizon may facilitate analytical methods that provide the decision maker insight. In many cases, the resources provided to the decision maker have an “expiration date” beyond which they are no longer available to be allocated toward alternatives. As this date approaches, the decision maker may (or may not) be willing to modify their decision criteria. Continuous decision problems require analytical techniques that can reflect this facet of the decision maker’s thinking.

Taken together then, these three elements lead to the definition of a continuous decision problem as one in which the decision maker, within a finite time horizon, expects to sequentially engage in more than one decision event en route to a final selection. At each decision epoch, the decision maker has the ability to make zero or more selections from known alternatives, or to defer a decision until a later date when the set of alternatives may have changed.

A practical approach taken to address these types of problems today is seen in the Air Force Materiel Command’s (AFMC) AFMC approach to selecting DP Development Planning (DP) projects. In broad terms, this process utilizes a value model to assign a value score to each proposed effort, and then uses a LP linear programming (LP) model to maximize the sum of value scores of selected projects within the funding constraints of the DP program and the manpower constraints of the Air Force Product Centers. PMJ Professional military judgment (PMJ) is then applied to the resulting list to capture any considerations not explicitly modeled.

While DP effort proposals are accepted throughout the year, there is a final call for proposals in the October/November timeframe. All new proposals are scored against the value model and any data issues are resolved in December, and the LP model is run in January. Although there is an out-of-cycle process for handling DP requests that do not conform to this timeline, these efforts are generally required to provide their own funding. By collecting proposals throughout the year and engaging in a single portfolio construction effort, AFMC can better optimize the allocation of scarce resources, both in terms of dollars and manpower. The drawback is that a promising proposal may languish for the better part of a year if it arrives out of cycle and cannot secure its own funding stream.

1.2 Problem Statement

Miles and von Winterfeldt described decision analysis as consisting of “models and tools to improve decision making.” [30] What we seek then are models and tools that provide an improved portfolio construction process which preserves, to the greatest extent possible, the rigor and analytical defensibility of a traditional, monolithic construction process while also providing the flexibility to consider alternatives that arrive over time. For the purposes of this research, the focus is exclusively on a single-stage selection activity. That is to say an alternative, once selected, is included in the portfolio in the form it was considered. This distinguishes the problem from multi-stage methods, which will be discussed later, in which an alternative is selected and then periodically re-evaluated for continuation, modification, or termination. Ideally the method will accommodate varying levels of information regarding the decision problem, ranging from a naive case where there is little experience with the distribution of attribute values or the arrival rate of alternatives, to more experienced cases where there is a sound historical record with which to characterize these distributions. Finally, methods to characterize the level of uncertainty associated with the method’s outputs are sought, as these will be of significant value to the decision maker.

At the conclusion of this research effort, the following contributions will have been made:

- A robust definition of the class of problems termed “continuous decision problems”
- The identification of one or more effective methods for providing analytic support to decision makers facing a continuous decision problem that meet a basic set of criteria
- A clear identification of the aforementioned criteria
- Characterization of the tradeoff between decision quality and temporal flexibility gained by utilizing these methods
- A set of guiding principles in formulating and modeling continuous decision problems, including the selection of distributions and the selection and treatment of parameter values

The remainder of this document is organized as follows:

- Chapter 2 provides a review of the literature and a summary of the current state of the art
- Chapter 3 contains a paper published in 2014 in *The International Journal of Multicriteria Decision Making* describing the Triage Method
- Chapter 4 contains a paper published in 2014 in *The International Journal of Multicriteria Decision Making* describing extensions to the Triage Method to consider resource constraints
- Chapter 5 contains a paper describing the effects of decision problem parameters on the effectiveness of the Triage extensions. As this dissertation was finalized, this paper was under review by *The International Journal of Multicriteria Decision Making*

- Chapter 6 contains a paper describing methods for eliciting parameters required for the use of the Triage extensions. In addition, the paper explores various methods for applying these parameters to continuous decision problems. As this dissertation was finalized, this paper was under review by *Decision Analysis*
- Chapter 7 summarizes findings and provides suggested avenues for further research

II. Background and Related Work

Our key focus here is in selecting items to form a portfolio of some type. While there is a significant gap in the state of the art regarding continuous decision problems, there is an expansive body of literature and tools for the traditional portfolio selection problem. Choosing a methodology and toolset from this population can be a complex endeavor in its own right, and may add to the time and budget constraints already faced. [16] As with almost all decisions, the choice of decision methodology itself involves trade-offs between the level of accuracy required and the time and resources available to achieve it. [74] The vast majority of the literature on this topic proceeds from the assumption that the set of alternatives is known to the decision maker. Duncan summarizes a number of approaches with the assumption that the set of alternatives is known being so fundamental as to never be explicitly acknowledged. [27] In identifying gaps, Duncan highlights a number of significant issues that complicate the process of selecting an optimal portfolio, including the multiobjective nature of portfolio decision problems, the presence of uncertainty in candidate project measures, the prospect of interdependence between portfolio elements, and the social difficulty in gaining consensus among multiple decision makers with varied focus and priorities. The possibility that the decision maker(s) may need to evaluate project proposals independently over time is never mentioned.

Henig describes a successful application of decision analysis in the selection of R&D R&D projects. [43] The focus however is exclusively on the construction of the objectives hierarchy and the attributes and measures used to evaluate projects. The author explicitly states that the decision maker “has a finite number of projects and a finite number of versions for each project. The feasible set of alternatives is all the possible combinations of all the projects at different investment levels.”

Similarly Henriksen focuses on identifying an appropriate set of attributes and constructing a process that can be used to rank projects. [44] The focus here is on developing a selection methodology that maintains analytical rigor without introducing so much complexity as to render the method unworkable outside of academic applications. Ranking though is an activity that is only meaningful if the set is known. Indeed the author states “The ‘number’ generated as a result of the evaluation process is only useful for comparing and ranking alternatives *within that set*,” (emphasis in original). This statement is made in the context of overcoming the sense of “researcher animosity” associated with the perception that the project is being “graded.” It is indicative however of again encountering the fundamental assumption that the entire set of alternatives is known.

Polyashuk [76] suggests an approach in which two criteria types are used in multiple criteria model: those that “are used to characterize both the entire portfolio and its individual elements,” and those that “are solely used to evaluate the portfolio as a whole but not its elements.” The first set is assumed to be composed of quantitative measures, and is given priority over the second set which may consist of more qualitative criteria. Chien expands on this approach, focusing from the start on methods for evaluating the portfolio as a whole rather than its component pieces, as “the combination of individually good projects unnecessarily constitutes the optimal portfolio.” [19] The author proposes a new taxonomy of portfolio attributes: independent portfolio attributes, interrelated portfolio attributes, and synergistic portfolio attributes. Key gaps highlighted include the potential for interactions between portfolio elements and the exponential growth in possible portfolio combinations as the number of candidates grow. The development of measurement scales for the newly identified classes of attributes is identified as a topic for further research, but there is no mention of the possibility that the entire alternative set may not be known.

Lim investigates a search problem where the decision maker evaluates a sequence of options with the aim of selecting a single “best” alternative. [64] In this construct, with

each alternative the decision maker encounters, they must decide whether to select the option and terminate the search, reject the option and continue the search, purchase more information about an attribute of the option, or terminate the search in favor of a status quo option. Thus the decision maker is faced with a dilemma both in how deeply they should search within a given option, as well as how broadly they should search across the available options. The authors provide a dynamic programming approach to this problem and offer a method for determining optimal policies. The potential for adapting this dynamic programming approach, or utilizing the similar techniques of goal programming [29] or preference programming [63], to the construction of a resource constrained portfolio of options rather than a single alternative is a potential avenue for further research.

Perhaps the most expansive body of related research in the literature deals with constructing portfolios that maximize financial gain. Examples include but are by no means limited to pharmaceutical R&D projects as in [84], stocks and other securities as in [61], or oil field explorations as in [13]. None of these approaches capture all the aspects of the continuous decision problem as described.

Many R&D evaluation methods recognize two types of decision: the original selection of a project for pursuit, and sequential decisions about whether or not to continue with the project based on information gained in the current stage of research. [1] This relates them to the sequential search problem in the sense that the decision is whether to continue to purchase additional information about the option or terminate consideration of that option. Bayesian approaches are particularly attractive in this setting, and may be extended to form what are known as influence diagrams [48] to provide a more complete modeling of the decision situation. Decision trees [15, 39, 40] are another widely used tool in this situation and providing a decision maker with an optimal chronological sequence of decisions. The solution this research pursues is more closely coupled with the first decision type, portfolio selection. While there may be applicability of the techniques associated with these methods

as mentioned previously, no explicit treatment of a temporally evolving set of alternatives is mentioned.

Exploration problems tend to assume again that the set of options is known and focus on methods for reducing the exponential growth of pairwise comparisons necessary to form a joint probability distribution. This distribution is then used to order (i.e. rank) the exploration opportunities.

There is an obvious analogy in selecting project portfolios to selecting portfolios of financial instruments, and it is useful to ask what techniques from this realm may apply to our problem. [1, 6, 19, 35, 61] One of the more established financial methodologies is NPV Net Present Value (NPV) analysis (sometimes also referred to as discounted cash flow analysis). NPV involves establishing a *discount rate* that represents the time value of money. That is to say, the discount rate establishes the decision maker's preference for a dollar today versus a dollar at some time in the future. The expected cash flow of a project can then be discounted by this rate, and the value *today* for its anticipated future performance can be established, and is known as the net present value. If all available alternatives are represented as such discounted cash flows, the decision maker can choose the set that maximizes total NPV. [9] A clear limitation to this approach is the need to monetize, as it can be very difficult to accurately predict the cash flows associated with a project proposal, and errors in forming this estimate can have a dramatic effect on the composition of the selected portfolio. The selection of the discount rate can also have a profound impact on the financial performance of an alternative, leading de Neufville to declare "The choice of the discount rate is the single most critical element in any evaluation of benefits and cost over time." [25]

A growing area of interest in applying financial methods deals with a class of techniques known as real options methods. [11, 80] In a real options framework, "any corporate decision to invest or divest in real assets is simply an option. Option holders

have the right but not the obligation to make an investment...” [70] There are limitations in seeking to use these methods for the portfolio construction problem under consideration. To begin with, most applications of real options, like NPV analysis, require the benefits of alternatives to be monetized, and suffer from all the drawbacks associated with this activity. Additionally, being an outgrowth of a financial sector characterized by open, efficient markets, real options may suffer when applied to military or governmental decisions. In these cases, certain assumptions may not hold, particularly regarding the openness of the market. For example, the government does not have market data to turn to in determining a reasonable price for an asset. In most cases, a government decision maker also does not have the option to simply do nothing and hold onto their money. On the other side of the transaction, producers are typically not free to market their products elsewhere if the government declines to purchase them.

The mathematical approach known as multiobjective optimization is among the most rigorous, but also the most difficult in practice to implement. It seeks to define multiple objective functions and minimize the value of those functions subject to one or more constraints. In practice, the objective functions are almost always contradictory such that a change that decreases the value of one of the objective functions will often increase the value of one or more of the others. The solutions to multiobjective optimization problems are those that are members of what is known as the Pareto optimal set. The Pareto optimal set is defined as the set of solutions consisting of “objective vectors such that none of the components of each of those vectors can be improved without deterioration to at least one of the other components of the vector.” Unfortunately, even for simple problems, the size of the Pareto set grows exponentially with the number of inputs. [90] Optimization algorithms may identify solutions that appear optimal, but may in fact correspond to local minima that are dominated by other solutions. Strict determination of whether a particular solution is

Pareto optimal is NP-hard [20] and often requires exhaustive enumeration of the Pareto optimal set.

The most directly applicable approach identified to date is the Triage Method described by Gutman. [38] Before describing the application of the Triage method, it should be noted that it requires a value model in order to be executed. The basic methodology for constructing and using a value model is described by Kirkwood. [59] The key elements of the value model are the value hierarchy, the weights applied to that hierarchy, and the evaluation measures. It is important to keep in mind three desirable properties of the value hierarchy. It should be *complete* in that it covers all the significant evaluation concerns that are associated with achieving the decision maker's objective. This property is sometimes referred to as being "collectively exhaustive". The hierarchy should be *non-redundant* in that no evaluation criteria should overlap. This property is sometimes referred to as being "mutually exclusive" and is important to ensure that aspects of the alternatives are not "double-counted". Finally, the hierarchy elements should display *preferential independence* in that the preference displayed for the achievement of one objective is not affected by the level achieved on any other objective.

Also key to the application of the Triage method is the concept of *sensitivity analysis*. Sensitivity analysis is a complex topic in its own right, and there are a number of methodologies for conducting it. [14, 50] Fundamentally, sensitivity analysis provides a method to gauge the impact of uncertainty in the weights assigned to the elements of the value hierarchy. Typically, the analyst will assist the decision maker in specifying bounds over which the weight of a hierarchy element may vary. The analysis then varies the weight assigned to that element over the possible range, varying the weight of the other hierarchy elements correspondingly such that all weights sum to 1, to determine the impact on the overall score of the alternative. The Triage Method uses global sensitivity analysis

to compare a single alternative's best and worst case performance potential to a cutoff value α . The method is applied as follows:

1. A value model for evaluating alternatives is developed, including a value hierarchy, weights, and single dimensional value functions. Sensitivity intervals for each of the weights in the value hierarchy are specified, and a cutoff value α is determined.
2. As an alternative A arrives, it is scored against the value model to determine a value score, $V(A)$.
3. Linear programming is used to determine the set of weights W_{max} that maximize $V(A)$, consistent with the previously specified intervals.
4. Simultaneously, a minimum value for $V(A)$ consistent with the specified weight intervals is determined
5. $V_{max}(A)$ and $V_{min}(A)$ are compared to α and the alternative is "triaged" into one of three categories:
 - (a) Where $V_{min}(A) > \alpha$ the alternative is selected
 - (b) Where $V_{max}(A) \leq \alpha$ the alternative is rejected
 - (c) Otherwise, the alternative is held pending further analysis

The Triage Method provides a fast, analytically defensible method for continuously evaluating alternatives as they arrive. Unlike methods applied to static decision situations, where the ranking of an alternative's value score relative to other alternatives is the focus, this method keys on the value of each individual alternative. Each alternative is allowed to reach its maximum potential value score within the weight space specified by the decision maker. Since the set of weights that maximize one alternative's value are likely different from those that maximize another's, it is not useful to compare the maximum value scores

to one another. Instead, they are compared to the cutoff value α . Gutman et al. also propose a value of α that changes over time [38]. This modification will be discussed in Chapter 4. A further open question is the best method for establishing the value of α , which is some ways comparable to the problem of selecting a discount rate for NPV analysis.

As initially proposed, the Triage Method is primarily used as a screening tool to rapidly evaluate a stream of alternatives and determine which are promising candidates for further, more detailed analysis. This is highly applicable in a case where the final decision is the selection of a portfolio of alternatives as opposed to a single alternative. A desirable extension to the Triage methodology would be to allow its use for portfolio selection in addition to screening.

The concept of triage is not new, and methodologies to accomplish similar functions have been proposed. For instance, Spradlin and Kutoloski [82] propose a construct where again alternatives are separated into one of three categories: the doomed projects, the equivocal projects, and the favorite projects. The similarity to our proposed selected/rejected/pending partition is obvious. Here again though, the implicit assumption is that the complete set of alternatives is available when the portfolio selection decision is being made.

III. The Triage Method

*The contents of this chapter were published in 2014 in The **International Journal of Multicriteria Decision Making**. They have been reformatted to comply with the AFIT style guide. The work is primarily that of Dr. Weir and Alex Gutman, with minor inputs from Jeremy Hendrix. It is included here for completeness as it forms the foundation for the remainder of the work presented in this dissertation.*

3.1 Introduction

“The theory of decision analysis,” as stated by [52], “is designed to help the individual make a choice among a set of prespecified alternatives.” Many researchers have examined how to choose the best alternative, or subset of alternatives, from this well-defined set; comparatively few, however, have analyzed how those alternatives initially arrived at the decision point. One of three scenarios is likely: the alternatives were generated by the decision maker, as in value-focused thinking [54]; they were specified in advance by an outside party; or they were screened from a larger pool of potential alternatives. We present the Triage Method to aid decision makers with the third scenario, in which a large set of alternatives must be screened to produce a smaller, more manageable number that can be thoroughly scrutinized before a final choice is made. Additionally, we consider the added caveat that alternatives may arrive over time and may expire.

Screening can reduce the time and cost of carrying out detailed analysis on undesired alternatives. For example, [77] discussed screening’s potential application in “site selection, new product decisions, executive recruitment and evaluation, evaluation of projects in education and health systems, and selection of corporate plans and strategies.” [53] carried out a meticulous screening of potential sites for an energy facility. [85] applied screening to reduce the set of policy analysis alternatives, and [4] applied a three-stage

screening method to projects providing fresh water to Newport News, Virginia, USA. Other applications include screening research and development proposals [83] and stocks [78].

In [18], a general screening procedure is defined as a function Scr with a set of alternatives, $\mathbf{A} = \{A^1, A^2, \dots, A^m\}$, such that

$$\emptyset \neq Scr(\mathbf{A}) \subseteq \mathbf{A},$$

where $Scr(\mathbf{A})$ denotes the remaining alternatives. The set of alternatives $\mathbf{A}$ is implicitly assumed to be a fixed set prior to applying the function Scr , and literature on multiobjective screening techniques and applications follows this paradigm. The screening techniques presented by [67] and [17] also assume the set of alternatives is fixed. This, however, does not encapsulate all decisions. Waiting for all alternatives to present themselves could hinder the screening process. It is not always practical to wait for one specific time to screen alternatives and another time to make a final decision, be it choosing a single alternative or constructing a portfolio. It can lessen the burden of the decision maker to quickly screen alternatives, remove those not capable of satisfying the decision objective, and thoroughly vet the transferred alternatives in batch sizes *at the same time*. In this paper, we relax the notion that $\mathbf{A}$ is fixed and seek to develop the triage method to expedite the screening process by evaluating alternatives as they arrive over time. The triage method is meant to aid in two-stage, ongoing decision processes which involve

1. screening alternatives
2. in-depth analysis of the alternatives that passed screening.

Consider the following example.

A company has one month to hire ten new computer experts to compete for a government contract. It posts the job announcement on its website and starts to receive resumes immediately. This two-stage decision process involves


1. screening resumes

2. interviewing candidates.

Given the month time-constraint, it's not prudent for the company to wait two weeks before screening resumes to select a group to interview. If the company selects 20 candidates to interview, what if more than ten interview poorly, embellish their resumes, or decide they do not want the job? The company would not want to be in the position of collecting more resumes with so little time before the deadline. Or, they may also find a highly qualified applicant, having received no response, has moved on to other opportunities and is no longer available to interview (the applicant has 'expired'). Realistically, the company would concurrently screen resumes as they arrive and interview qualified applicants throughout the entire month. Applicants would be hired one-at-a-time or in batch sizes until the decision objective was met (i.e., hire ten qualified computer experts to support the company's contract bid). Additionally, the person screening resumes may decide - as the deadline approaches and hiring is not complete - to lower the standards required for an interview. This would give an opportunity for applicants on a wait list or those still submitting resumes to interview for the open positions. Or, it may be the case that the interviewer is seeing too many applicants, so they would require a higher standard for an interview.

Many situations are analogous: suppose you are looking for a new house or car; you will most likely screen alternatives online while visiting houses or test-driving cars that passed your screening criteria. An academic journal screens articles while pushing some forward to be officially reviewed. Scholarship applicants must meet certain criteria before being interviewed. A college admissions board cannot wait until all applications are submitted before making acceptance decisions because the best prospects may choose other schools if they do not hear back from the admissions board. Difficult decisions are being made in situations like this, but the current scope of screening does not reach this type of decision.

Rather than a fixed set of alternatives, $\mathbf{A}$, we consider a changing set of alternatives $\mathbf{A}_t$ with a time-dependent screening function, $Tr(\mathbf{A}_t, \alpha_t)$, called the triage function, which partitions alternatives into three groups via multiobjective decision analysis: transferred (i.e. move to stage two), rejected, and pending. Alternatives are partitioned based on their value relative to a changing cutoff value, α_t . The process takes place over time, t . Figure 3.1 depicts the Triage Method at time t . The set of alternatives $\mathbf{A}_t$ is comprised of $\mathbf{N}_t$, the new, unknown alternatives entering the decision process at time t , and $\mathbf{P}_{t-1}$, the pending alternatives from $Tr(\mathbf{A}_{t-1}, \alpha_{t-1})$. Because the triage function changes over time, the pending alternatives are cycled back through the process, as they could eventually be transferred.



We present the triage method to deal with the four following screening scenarios:

1. *Continuous decisions with no final time limit:* For example, a research and development team must manage a portfolio where alternatives are continuously added over time, but the portfolio is never ‘fixed’ i.e., at no time, t , will the decision maker say “This is our portfolio. We are done”.
2. *Decisions that must be made quickly because the alternatives can expire:* For example, if someone is looking for a new house, he or she will screen houses to

visit and eventually buy one. The time factor here is not necessarily a deadline on buyer's part. The buyer may be able to live where they are now indefinitely. However, time is a factor because their favorite house may be sold. They buyer must choose between buying the house or risk losing it if they wait for something better.

3. *Decisions that have a time limit and alternatives can expire*: Same example as above, but suppose the buyer has a deadline on the decision because their apartment lease expires at the end of the month.
4. *Continuous decisions with a time limit*: For example, a hiring manager must bring in people to interview and fill positions. For a government contractor, there is a deadline on the decision maker to have his/her portfolio complete by a fixed decision point, t . There is the added constraint that the hiring manager must make decisions quickly because the top recruits may leave for somewhere else. This is a combination of points 2 and 3.

All four scenarios share a key commonality. There is never a point where the decision maker can wait for all alternatives to present themselves before making a decision. While it would be possible to rank alternatives at different time points, it would be more beneficial to immediately screen out the unwanted alternatives to save time. In each situation, we reject poor alternatives from the set A_t while giving a more detailed inspection to those that passed screening. It is important at this stage to note the two-stage construction of the problem and to make clear that the triage method is designed to address the first stage. It is a tool for interim decisions. The second stage selection of one or more alternatives from a well-defined set is a well-established problem that can be approached with any number of methods found in the literature (Keeney, 1992; Yoon and Hwang, 1995; Keisler, 2004). As noted by a reviewer, the proposed method is also similar to other approaches in the literature. For example, the ELECTRE TRI method (Mousseau and Slowinski, 1998)

offers ways to choose, rank, and sort actions based on preference, and the UTADIS method (Zopounidis and Doumpos, 1999) uses cutoff values based on the inference of additive value functions. What these methods require is a fixed set of alternatives from which to choose. The triage method provides a fast, rigorous, defensible process for screening a large number of potential alternatives in order to provide this necessary initial condition. Additionally, it does so in a manner that allows alternatives to arrive asynchronously over a period of time. In the next section, we review current screening techniques and outline the motivation for the triage method. In Section 3, we formally describe the triage method in detail, and in Section 4, we illustrate a wartime application of the triage method developed for the Joint Improvised Explosive Device Defeat Organization (JIEDDO). Last, we summarize results and discuss the benefits of the triage method.

3.2 Principles of the triage method

Screening is usually accomplished by eliminating poor alternatives that do not satisfy or exceed minimum criteria. [89] discuss two such “satisficing” methods: Conjunctive and Disjunctive. [34] mathematically expressed these methods along with Compensatory screening. Each method requires the decision maker to establish cutoff values. With conjunctive and disjunctive screening, alternatives are screened into a ‘choice set’ using an indicator function $I(x_j, \gamma)$ that equals one if the screening criteria, x_j , is greater than the cutoff value, γ ; the function is zero otherwise. However, using indicator functions on an alternative’s attributes will not consider trade-offs. “Using screening criteria to imply value judgments”, according to [55], is a common mistake in making value trade-offs, and failure to consider trade-offs during screening could eliminate potentially good alternatives. Compensatory screening is preferable because it considers trade-offs.

In his screening model of energy facility sites, [53] describes compensatory screening as a “decision analysis model” because it allows value trade-offs; an alternative can have some weak attributes as long as it compensates the lost value with excellent attributes.

The difficulty with compensatory screening lies in the selection of the cutoff level, which takes good judgment by subject matter experts. [53] also describes a comparison screening method which considers value trade-offs between a fixed set of alternatives. For the Triage Method, we combine *compensatory* and *comparison* screening.

For the compensatory screening aspect of the Triage Method, we exploit the usability and transparency of value hierarchies to structure objectives and consider trade-offs. [52] explain the need for a hierarchy's operability, decomposability, nonredundancy, and small size. Weights, denoted w_i , are assigned to each objective to specify the value trade-offs of one objective to another. An alternative's value $v(A^j)$ with respect to n attributes $\{a_1^j, a_2^j, \dots, a_n^j\}$ is aggregated with the function

$$v(A^j) = \sum_{i=1}^n w_i v_i(a_i^j) \quad (3.1)$$

where $v_i(a_i^j) \in [0, 1]$ is the scaled rating of a_i^j . The weights are normalized, so an alternative's overall value ranges from 0 to 1. See [59] for more information on how to properly construct and weight value hierarchies. [53] and [85] stress the need for a simplified model with readily available data when screening. Further, the attribute levels, a_i^j , should have little or no uncertainty. Collecting attribute data with significant uncertainty can slow the screening process and cause disagreements between the involved parties.

Since the set of alternatives $\mathbf{A}_t$ changes over time, we cannot perform a complete comparison of all alternatives relative to one another, as done by [53]. Rather, for the comparison aspect of the triage method, we consider hypothetical decision alternatives. For example, consider a desirable, hypothetical decision alternative $D = \{d_1, d_2, \dots, d_n\}$ where each d_i represents a cutoff level for the i th objective. Decision makers and managers should have a good idea of what makes a desirable cutoff for each of the simplified attributes. Using the value model in Equation 3.1, we can calculate a cutoff level $\alpha = v(D)$ and compare the value of incoming alternatives, which arrive sporadically, to α . Value trade-

offs are inherent in the hierarchy, so an alternative A^j can satisfy $v(A^j) \geq \alpha$ even though $v_i(a_i^j) < v_i(d_i)$ for some i . This concept is explored further in Section 3.3.3.

3.2.1 Decisions over time.

Theoretical decision analysis models often discount the role of screening, but research suggests screening plays an integral, underappreciated role in decision making. [10] discusses the important aspects of behavioral decision making that are generally overlooked during screening, and most important to this research is the element of time. Peoples' preferences change over time, and when screening alternatives, decision makers may search or wait for more alternatives to present themselves if the current set is undesirable. For the Triage Method, we assume the alternatives are unknown and submitted from an outside source, so if screening does not transfer enough alternatives to satisfy the decision objective, it's necessary to rescreen the set of pending alternatives. [10] examined how decision makers react when the set of transferred alternative is empty and concluded two important results: "subjects lowered their standards and became more tolerant of violations of those standards."

It may also be the case that the decision maker raises their standards as time goes on. Whatever the decision scenario, the decisions in stage two will have an impact on the screening in stage one, particularly with respect to the cutoff value. As time goes on, should the decision maker increase or decrease the cutoff value? The second stage decision plays a role in assessing this, and though we do not formalize the second decision stage in this article, we provide guidance on using and adjusting the cutoff value in later sections.

3.2.2 Cost considerations.

Because the triage method is not a choice model, we recommend that users do not include cost, as in cost-benefit analysis. The triage method should help decision makers partition alternatives in the value space. Whether an organization can or will ultimately fund a project, either individually or as part of a portfolio, is a choice that will include

factors such as cost not in the screening model. We feel this is also consistent with ensuring as little uncertainty in the data as possible. Certainty in cost about future alternatives usually requires in-depth analysis and does not lend itself to the quick analysis we propose in our method.

Moreover, our focus is on the first stage of this two-stage decision problem. Cost is likely a concern for many decision problems, but our intent here is not to focus the conversation on cost because it can dictate the entire discussion. No matter how inexpensive some alternatives are, they will not be chosen at the final decision point if they have little value. While a low cost may inflate an alternatives cost-benefit value above a higher-valued, but costlier alternative, this first stage is to focus on the value of an alternative. There may be a risk that those alternatives sent to stage 2 are prohibitively expensive. This can be avoided by having a cap for the cost of an individual alternative. For example, a decision maker may not want to allocate more than 110th of hisher resources on one alternative. Also, not all decisions are heavily weighted on cost. Cost is not an issue when screening college applicants or screening applicants for a job with fixed pay.

3.3 The triage method

Screening will stop analysis of unwanted alternatives, and the triage method will accomplish this quickly. Again, the triage method is for the first stage of a two stage decision process. A different model would be used in stage two of the decision process to select the final alternative or portfolio of alternatives to satisfy the objective. Researchers have applied similar partitioning techniques for a final choice model. For instance, [81] discuss a triage rule which partitions research and development projects into “the doomed projects, the equivocal projects, and the favorite projects”. [57] expanded their simplified process using a distance from a threshold as a Triage Rule to automatically fund projects into a portfolio. However, they focused on proceeds per unit of funding, which is similar

to a cost-benefit analysis. Our triage method, however, focuses on screening over time and consists of five components, which will we discuss in this section:

1. a value model
2. hypothetical decision alternatives
3. weight intervals
4. cutoff values
5. comparison strategy: pessimistic or optimistic approach

We assume a simplified value hierarchy with appropriate functions $v_i(\cdot)$ has been developed and weighted to accurately define value trade-offs among different objectives. To be consistent with decision analysis screening models, attribute data for the incoming alternatives should be easy to capture and have little or no uncertainty [53]. Table 3.1 provides a summary of the notation and concepts used in the Triage Method.

3.3.1 Hypothetical decision alternatives.

As alluded to in Section 3.2, we apply a comparison aspect to the triage method. When screening alternatives, a decision maker should consider their needs, wants, and desires (West, 2011). Early in the screening process, decision makers will want to prioritize alternatives with attributes exceeding the minimal, needed level of attainment. Therefore, for each objective, the decision maker should specify a desirable attribute which provides a ‘margin of excellence’ to the given alternative (West, 2011). This constructs a hypothetical ‘Desired’ alternative, D with attributes $\{d_1, d_2, \dots, d_n\}$ which creates a higher, prioritizing cutoff value for incoming alternatives.

“Needs,” much like the cutoffs γ_m in the conjunctive method, specify an essential level of attainment for a given attribute. Decision makers or stakeholders can think of needs as “must have requirements” [87], so for each objective, a needed attribute, n_i ,

Table 3.1: The Triage Method Notations and Descriptions

Notation	Description
T	Time constraint to complete screening
t	Current time period: $t = 0, 1, 2, \dots, T$
D	Hypothetical “desired” alternative with measures $\{d_1, d_2, \dots, d_n\}$
N	Hypothetical “needed” alternative with measures $N = \{n_1, n_2, \dots, n_n\}$
W_S	Weight space to capture uncertainties in original weights
α_0	Highest cutoff value: $\alpha_0 = \underline{v}(D)$
α_T	Lowest cutoff value: $\alpha_T = \underline{v}(N)$
$f(t)$	Monotonically decreasing time function with $f(0) = 1$ and $f(T) = 0$
α_t	Cutoff value at time t : $\alpha_T + f(t)(\alpha_0 - \alpha_T)$
$Tr(\mathbf{A}_t, \alpha_t)$	Triage Function with inputs $\mathbf{A}_t$ and t

should be chosen. This constructs a hypothetical “Needed” alternative, N , with attributes $\{n_1, n_2, \dots, n_n\}$ which creates the lowest possible cutoff value for incoming alternatives. With value trade-offs, it is possible for an alternative to have an attribute level below n_i , but the Triage Method will not immediately reject the alternative if its *overall* value is above the aggregate cutoff level, $v(N)$. By construction, $v_i(n_i) \leq v_i(d_i)$ for all $i = 1, 2, \dots, n$. The values $v(D)$ and $v(N)$ will guide our choice of a cutoff level.

It is important to realize the difference between

1. the ‘needed’ and ‘desired’ attribute levels
2. the minimum and maximum attribute levels that are scored by the value functions, $v_i(\cdot)$

A decision maker would obviously prefer an alternative whose attributes attain the maximum value for each objective, but that does not mean the “desired” alternative D should satisfy $v(D) = 1$. It is not likely an incoming alternative will have the highest

possible values for all objectives so setting $v(D) = 1$ would reduce the practicality of the model and equate more to the conjunctive rule, which does not specify trade-offs. Suppose, for example, a hiring manager was screening job candidates based on education level and years' experience. The minimum level for education could be high school diploma/GED (value of 0) and the maximum level could be a doctorate degree (value of 1). While a doctoral degree has more value than a Master's, the "desired" level could be a Master's degree and the "needed" level could be a bachelor degree. Anything above desired attribute levels will most likely be countered with attributes levels below a different attributes' desired or needed level.

3.3.2 *Accounting for uncertainty with weight intervals.*

Selecting or rejecting an alternative based solely on a comparison to α_i could be scrutinized as a hard-line approach to screening. The value of an alternative may be within a small margin of α_i and the choice to transfer or reject may not be clear cut. While the attribute levels should be clear in screening, any value model has intrinsic uncertainty because of the weights, $w_i, i = 1, 2, \dots, n$. Weights are subjective, and decision makers may not completely trust the results of the model if they cannot accurately specify the relative importance and trade-off between competing objectives [31].

The most direct method for dealing with uncertainty in the weights is to perform sensitivity analysis. This analysis varies the weights within pre-defined bounds and determines the changes in the resulting value scores. [60] describes three such interval bounds:

1. $w_i^* \in [w_i \pm \lambda_i]$ (strict intervals)
2. $w_i^* \in [1 \pm \lambda'_i]w_i$ (relative intervals)
3. $w_i^* \in [w_i(\min), w_i(\max)]$

where w_i^* represents the new weight of the i th objective and λ_i and λ'_i are the maximum allowable change or relative change from the original weight w_i , respectively. [60] set $w_i(\min) = \max\{(w_i - \lambda_i), (1 - \lambda'_i)w_i\}$, and $w_i(\max) = \min\{(w_i + \lambda_i), (1 + \lambda'_i)w_i\}$ to combine strict and relative intervals so the smaller weights use relative intervals and larger weights use strict. The new weights must satisfy $\sum_{i=1}^n w_i^* = 1$, and linear programming is used to accomplish this.

The weight space effectively produces a range of values for each alternative. Let $\bar{v}(A)$ and $\underline{v}(A)$ denote the maximum and minimum potential values of $v(A)$, respectively. Define the spectrum of potential values for alternative A , denoted $S(A)$, as the range

$$S(A) = [\underline{v}(A), \bar{v}(A)]. \quad (3.2)$$

$S(A)$, therefore, summarizes A 's range of values subject to the decision maker's preferences and value trade-offs. Further, $S(A)$ encapsulates the alternatives strengths and weaknesses.

For an alternative A , linear programming can be used to find the weight vectors $\overrightarrow{w_{max}}$, $\overrightarrow{w_{min}}$ in a weight space that maximize and minimize A 's overall value. If the weight space, W_S , is defined with strict or relative intervals, the benefit of the Triage Method is the instantaneous calculation of $\underline{v}(A)$, and hence $S(A)$, after linear programming evaluates $\bar{v}(A)$, as shown below.

Let A be an alternative with n attributes $\{a_1, a_2, \dots, a_n\}$ and original weight vector $\overrightarrow{w_O} = \{w_1, w_2, \dots, w_n\}$. The overall value of A , $v(A)$, is defined in Equation 3.1. We wish to maximize the linear function $v(A)$ by altering the weight vector within the specified weight space, W_S . Without loss of generality, assume W_S was defined with relative intervals. Note that λ_i can be different for each weight in W_S . For each weight w_i , we can alter the $v(A)$ by choosing $s_i \in [-\lambda_i w_i, \lambda_i w_i]$ to change the weight of attribute a_i from w_i to $w_i + s_i$. Since $\sum_{i=1}^n (w_i + s_i)$ must equal one, it follows that $\sum_{i=1}^n s_i = 0$.

Let $S^* = \{s_1^*, s_2^*, \dots, s_n^*\}$, such that $s_i^* \in [-\lambda_i w_i, \lambda_i w_i]$ and $\sum_{i=1}^n s_i^* = 0$, be the set that maximizes $v(A)$. i.e. $\overrightarrow{w_{max}} = (w_1 + s_1^*, w_2 + s_2^*, \dots, w_n + s_n^*)$. We are guaranteed such a set

exists because $v(A)$ is a linear function. Thus,

$$\bar{v}(A) := \max_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i + s_i^*) v_i(a_i) \quad (3.3)$$

The calculation of $\min_{\vec{w} \in W_S} (v(A))$ is immediate by the following theorem.

Theorem 1. *Let $S^* = \{s_1^*, s_2^*, \dots, s_n^*\}$ with $s_i^* \in [-\lambda_i w_i, \lambda_i w_i]$ and $\sum_{i=1}^n s_i^* = 0$ be the set that maximizes $v(A)$. i.e.*

$$\max_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i + s_i^*) v_i(a_i).$$

Then,

$$\underline{v}(A) := \min_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i - s_i^*) v_i(a_i). \quad (3.4)$$

Equivalently, $\vec{w}_{min} = (w_1 - s_1^*, w_2 - s_2^*, \dots, w_n - s_n^*)$.

See the appendix for the proof. A similar proof validates the case with strict intervals. Theorem 1 verifies $v(A)$, defined by Equation 3.1, is the midpoint of $\bar{v}(A)$ and $\underline{v}(A)$. Thus, the calculation of $\underline{v}(A)$ is simply $\underline{v}(A) = 2v(A) - \bar{v}(A)$.

The decision maker can have more confidence in a value that remains high over the range of plausible weights than in one that drops significantly. The Triage Method exploits this technique to highlight an alternative's performance. At first glance, this may appear to provide an avenue to manipulate the model to produce a pre-determined result or favor a preferred alternative. Were it the case that weights and sensitivity intervals were defined after the alternatives had been considered, this might be the case. It is important to note that the model is constructed, weights are assigned, and sensitivity intervals are determined *prior* to the consideration of any alternatives. Indeed, in the situations the triage method was developed to support, the alternatives are likely not even known at the time the model is constructed. Once the alternatives begin to reveal themselves, the only parameter open to adjustment is the cutoff value, and changes to this parameter potentially impact all alternatives.

3.3.3 Cutoff value and comparison strategy.

Next, we need to assign a cutoff value α_t to compare each alternative's spectrum of values. The uncertainty in the weights and calculation of $S(A^j)$ for each alternative necessitates the calculation of $\underline{v}(D)$ and $\underline{v}(N)$, the lowest possible values for the “Desired” and “Needed” alternative, respectively. At all times, t , α_t must be in the range defined by $[\underline{v}(N), \underline{v}(D)]$. The upper bound is set to $\underline{v}(D)$ rather than $\bar{v}(D)$ to protect an alternative A^j satisfying $a_i^j = d_i$ for all $i = 1, 2, \dots, n$ from entering $\mathbf{P}_t$ because $\underline{v}(A^j) \not\geq \alpha_0 = \bar{v}(D)$. The initial value, α_0 , should be chosen to reflect the decision maker's expectations for what alternatives should initially pass screening. Then, as time changes, the decision maker can adjust the value of α_t to a new value in $[\underline{v}(N), \underline{v}(D)]$.

Comparing $S(A)$ to the cut-off score α_t gives the decision maker substantial justification when screening alternatives. Further, the weights maximizing one alternative's value will not necessarily maximize another's, so this technique gives each alternative the opportunity to reach its absolute potential. The decision maker can then triage alternatives into three groups - transferred (**T**), rejected (**R**), or pending (**P**) - using the Triage Function:

$$Tr(\mathbf{A}_t, \alpha_t) = \begin{cases} \mathbf{T} = \{A^j : \underline{v}(A^j) \geq \alpha_t\} \\ \mathbf{R} = \{A^j : \bar{v}(A^j) < \alpha_T\} \\ \mathbf{P}_t, \text{ otherwise} \end{cases} \quad (3.5)$$

If $\bar{v}(A^j) < \alpha_T$, then A^j should be rejected because even under optimal conditions, it falls below the needed value. As written in Equation 3.5, if $\underline{v}(A) \geq \alpha_t$, then A^j should be transferred because its value under a worst-case scenario still surpasses the desired cutoff level. (See [73] for benefits of using a worst-case analysis.) If $\alpha_t \in S(A)$, A^j is classified as a pending alternative and will be cycled back through the process.

Using a worst-case value to screen alternatives into **T** is a pessimistic approach because alternatives are transferred only if their worst possible values surpass the cutoff.

This is similar to the maximin technique in [89]. Should the situation dictate an optimistic approach (a variation of the maximax technique), the Triage Function could transfer all alternatives that satisfy $\bar{v}(A^j) \geq \alpha_t$ because there is a *possibility* the alternative's true value is above the cutoff. Therefore, $\mathbf{T}$ in Equation 3.5 could be defined as:

1. Pessimistic approach: $\mathbf{T} = \{A^j : \underline{v}(A^j) \geq \alpha_t\}$
2. Optimistic approach: $\mathbf{T} = \{A^j : \bar{v}(A^j) \geq \alpha_t\}$

Assuming equivalent alternatives, the Pessimistic approach will transfer fewer than the Optimistic approach. Using the Optimistic approach for the Triage Method will completely partition the alternatives at $t = T$ since either $\bar{v}(A^j) \geq \alpha_T$ (i.e. A^j is transferred), or $\bar{v}(A^j) < \alpha_T$ (i.e. A^j is rejected). The Pessimistic approach, however, will most likely have a nonempty set of Pending alternatives when $t = T$. The Triage Function only rejects the alternatives that, at their best value, do not surpass the needed cutoff. Thus, it is possible that $\alpha_T \in S(A^j)$, meaning A^j will never satisfy $\underline{v}(A^j) \geq \alpha_t$. The resulting pool of Pending alternatives, $\mathbf{P}_T$, could easily be reconsidered if the situation required it. Both the pessimistic and optimistic approaches are acceptable and dependent on the decision situation. The decision maker must decide if it is important to reduce the number of transferred alternatives, which could possibly eliminate a good alternative from contention, or if the decision maker is risk averse and has time to review more alternatives, they could apply the optimistic approach. Sending more alternatives to stage two may end the decision at an earlier time, so there is a risk of missing a great alternative that arrives after the final decision point. There are advantages and disadvantages to each approach.

3.4 Illustrative example

This section applies the Triage Method to a proposal screening problem for the JIEDDO Joint Improvised Explosive Device Defeat Organization (JIEDDO). A team from the Air Force Institute of Technology created the value hierarchy [24]. At the present time,

the Triage Method has not been implemented at JIEDDO, as they recently reorganized shortly after we submitted the method. However, this section clarifies the theoretical components of our methodology and demonstrates its effectiveness.

3.4.1 Brief History of JIEDDO.

Roadside bombs and other homemade explosive devices pose a serious and deadly threat to coalition forces in Iraq and Afghanistan. Broadly known as IED improvised explosive devices (IEDs), these weapons have accounted for 70% percent of all U.S. combat casualties in Iraq and 50% in Afghanistan, killed and wounded, from 2003 - 2007 [88]. To synchronize counter-IED efforts, JIEDDO was established in 2006 as a permanently-manned entity. Its mission is to lead “Department of Defense actions to rapidly provide counter-IED capabilities in support of the combatant commanders and to enable the defeat of the IED as a weapon of strategic influence” [49]. The urgency of its mission is reflected in its budget; \$3.465 billion was appropriated for JIEDDO in fiscal year 2011 [23].

Compared to 2003, the weekly number of IED incidents in Iraq has dropped. However, coalition forces in Afghanistan have seen IED usage increase dramatically. During 2010, IED-related casualties in Afghanistan, including coalition forces and Afghanistan security forces and civilians, increased 19% [49]. This persistent violence waged against military and civilian targets via IEDs requires JIEDDO to aggressively find, develop, and deploy counter-IED capabilities to the warfighter by soliciting and funding counter-IED proposals from the military, academia, and industry [47].

Through BAA Broad Area Announcements (BAAs), JIEDDO communicates its countermeasure needs to outside organizations. Their acquisition goal is to acquire and deliver counter-IED initiatives to the warfighter within four to 24 months, so proposals are constantly submitted in response to BAAs. A panel of evaluators reviews batch sizes of proposals and partitions them into two groups: those with potential to defeat IEDs and

those with poor or infeasible concepts. A selected proposal then enters the JCAAMP Joint IED Defeat Capability Approval and Acquisition Management Process (JCAAMP), where it is more meticulously evaluated by a team of experts. The final decision to fund a proposal is made by the Deputy Director of JIEDDO [24].

JCAAMP was established in 2007 after “some in industry criticized JIEDDO for its *ad hoc* acquisition process and its inability to quickly and thoroughly evaluate proposals and provide feedback to industry” [26]. While the process shortened the time between development and deployment of counter-IED initiatives, the Triage Method has potential to further expedite this process by screening alternatives before they enter JCAAMP. From 2006-2007, 1,274 proposals were received by JIEDDO in response to the BAAs, and 447 passed initial review of the BAA Information Delivery System and entered the JCAAMP process [24]. The Triage Method can provide an understandable framework that would add credibility to the decision of transferring proposals into JCAAMP where a more thorough analysis takes place over several different stages.


3.4.2 Triage Method Applied to JIEDDO Model.

To thwart the threat of IEDs, JIEDDO established three Lines of Operation: Attack the Network (i.e. the terrorist teams that fund and create IEDs), Defeat the Device, and Train the Force. The overall objective, of course, is to defeat IEDs by any method. Accordingly, [24] specified the overall value on their hierarchy as “Potential to Defeat IEDs.” A brief summary of their hierarchy is provided here, but the crux of this section is the application of the Triage Method. For complete details on the hierarchy, including value functions and detailed descriptions, please see [24].

Figure 3.2 displays a representative value hierarchy developed for JIEDDO, and the ovals contain the performance measures for each item in the hierarchy. The reader will notice that cost is not included in the hierarchy, and this is because the purpose of this example is to identify proposals with promising counter-IED capabilities. Further, the

scope of the research is not limited to a specific Line of Operation because each could have its own funding requirements. Rather, the hierarchy encompasses the organization's mission at large. Cost would become a factor at later stages in the decision process, but the goal here is to screen out poor proposals that, regardless of funding, would not be acquired.

Figure 3.2: The Triage Method Hierarchy



[24] originally scored a set of 30 proposals that were previously evaluated by JCAAMP (i.e. they all passed the initial screening from the BAA Information Delivery System), and “the final breakdown included thirteen accepted proposals,...and seventeen rejected proposals covering all areas of submission for BAA.” Each proposal was given a proposal key because of classification issues. The rejected proposals’ keys are followed

by a # symbol. Their initial research concluded the hierarchy works well for prototype systems, so here, we analyze the 26 prototype proposals with the Triage Method.

The weights in this hierarchy were assigned with the swing weighting technique described in Section 4.4 of [59], and the weight intervals were assigned with relative intervals of $\pm 25\%$ (i.e. $\lambda_i = .25 \forall i = 1, 2, \dots, 13$) to demonstrate the technique. Again, the decision maker will specify the actual weight interval parameters, which can be different for each weight. Table 3.2 lists the hierarchy's weights, interval parameters, and weight space used for this example.

Table 3.2: Sample JIEDDO Hierarchy Weights and Weight Space

Value Hierarchy	Global Weights	Interval Parameter	Weight Space
Potential to Defeat IEDs			
Needed Capability			
Tenets Impacted	.056	25%	[.042, .070]
Gap Impact	.176	25%	[.132, .220]
Classification	.056	25%	[.042, .070]
Time to Counter	.112	25%	[.084, .140]
Operational Performance			
Technical Performance	.110	25%	[.083, .138]
Suitability	.056	25%	[.042, .070]
Interoperability	.091	25%	[.068, .114]
Technical Risk	.037	25%	[.028, .046]
Fielding Timeline	.056	25%	[.042, .070]
Usability			
Operations Burden	.087	25%	[.065, .109]
Work Load	.100	25%	[.075, .125]
Required Training			
Training Time	.050	25%	[.038, .063]
Program Maturity	.013	25%	[.010, .016]

As a proof of concept, we assume a six-month long submission process where proposals arrive in batches each month. The cutoff value will be adjusted monthly, so $T = 6$. Needed attribute levels, $\{n_1, n_2, \dots, n_{13}\}$ were chosen to construct a hypothetical “needed” alternative, N . Similarly, desired attribute levels $\{d_1, d_2, \dots, d_{13}\}$ were chosen to construct a hypothetical “desired” alternative, D . The monotonically decreasing time function $f(t)$ was chosen to be a simple linear function. The values of $v(N)$ and $v(D)$ were calculated using Equation 3.1 and then minimized using the procedures in the previous section with the weight space, W_S , defined in Table 3.2. Further, we assume a pessimistic approach of the Triage Method, using the Triage Function defined in Equation 3.5.


Given our sample of proposals, we will assume six proposals are present at the initial screening, $t = 0$, and ten proposals will arrive on $t = 1$ and $t = 2$. Thus, our sample of proposals represents only two months of the six-month screening process. Table 3.3 follows the flow of each proposal from the time they enter the process until screening is complete at the end of the six months. More proposals would enter the process at times $t = 3, 4, 5, 6$, but the purpose of Table 3.3 is to follow the proposals in our sample. The time-adjusted cutoff values also appear in the table. Figure 3.3 graphically displays the Triage Method results for time periods $t = 0, 1, 2$.

Based on the Triage Method, 2 proposals will be rejected, 17 will be transferred into JCAAMP, and 6 will remain in the pending pool of alternatives. The recommendations from the Triage Method are mostly consistent with the actual results from JIEDDO, and it’s implementation would reduce their workload of further vetting unwanted proposals. Specifically, if JIEDDO only analyzed the proposals transferred from the Triage Method, they would reduce their workload by about 35% because rather than moving all 26 into JCAAMP, only 17 would be transferred. Further, 12 out of the 13 funded proposals would be transferred into JCAAMP, so the recommendations from the Triage Method are consistent with future choice scenarios. The one anomaly is Proposal F, which received

Table 3.3: Triage Method Results on JIEDDO Proposals as t increases from 0 to 6.

Proposal	Potential	Minimum	$\alpha_0 = \underline{v}(N)$	α_1	α_2	α_3	α_4	α_5	$\alpha_6 = \underline{v}(D)$
ID	Value	Value	0.6739	0.6321	0.5902	0.5484	0.5066	0.4647	0.4229
A#	0.5226	0.3182	P₀	P₁	P₂	P₃	P₄	P₅	P
B	0.7169	0.6183	P₀	P₁	T	T	T	T	T
C	0.6194	0.5008	P₀	P₁	P₂	P₃	P₄	T	T
D#	0.4385	0.2854	P₀	P₁	P₂	P₃	P₄	P₅	P
E	0.6197	0.5016	P₀	P₁	P₂	P₃	P₄	T	T
F	0.4228	0.3064	R	R	R	R	R	R	R
G#	0.5626	0.4034		P₁	P₂	P₃	P₄	P₅	P
H	0.6835	0.5426		P₁	P₂	P₃	T	T	T
I	0.6164	0.4944		P₁	P₂	P₃	P₄	T	T
J#	0.6007	0.4777		P₁	P₂	P₃	P₄	T	T
K#	0.4680	0.3153		P₁	P₂	P₃	P₄	P₅	P
L#	0.4494	0.3219		P₁	P₂	P₃	P₄	P₅	P
M	0.6066	0.4824		P₁	P₂	P₃	P₄	T	T
N#	0.5245	0.3631		P₁	P₂	P₃	P₄	P₅	P
O	0.6870	0.5540		P₁	P₂	T	T	T	T
P#	0.6121	0.4927		P₁	P₂	P₃	P₄	T	T
Q	0.8667	0.7778			T	T	T	T	T
R#	0.5542	0.4304			P₂	P₃	P₄	P₅	T
S#	0.2104	0.1262			R	R	R	R	R
T	0.6172	0.4617			P₂	P₃	P₄	P₅	T
U	0.7873	0.6662			T	T	T	T	T
V#	0.6044	0.4520			P₂	P₃	P₄	P₅	T
W	0.6308	0.4978			P₂	P₃	P₄	T	T
X#	0.5474	0.4255			P₂	P₃	P₄	P₅	T
Y#	0.5501	0.3934			P₂	P₃	P₄	P₅	P
Z	0.6378	0.4985			P₂	P₃	P₄	T	T

Figure 3.3: Triage Method Results on JIEDDO Proposals



Each proposal's spectrum of potential values, $S(\cdot)$, defined in Equation 3.2 is represented above in comparison with the cutoff values at times $t = 0, 1, 2$. The center square of each proposal's spectrum represents its original value.

funding from JIEDDO but was rejected by the Triage Method. Assuming the hierarchy accurately reflects JIEDDO's objectives, Proposal F may have been a poor proposal for funding, and it may represent wasted resources.

3.5 Summary

The triage method is a time-dependent multiobjective screening process which assumes no knowledge of incoming alternatives. As alternatives arrive, they are compared to a hypothetical alternative to specify trade-offs and are then partitioned into one of three groups (transferred, rejected, and pending) via a triage function. The pending alternatives are cycled back through the process. It screens a continuous stream of an unlimited number of alternatives independently of one another to avoid delaying further analysis. The methodology is practical, transparent, repeatable, and computationally feasible for

hundreds of alternatives. Further, the model can specify why alternatives were rejected, providing justification and documentation to interested parties. The triage method was presented to guide decision makers when a large amount of alternatives must be reduced to a smaller, manageable size for further analysis. Literature on screening is comparatively small to literature on choice models, but screening can play a significant role in the decision process. Further, while current screening methods assume a fixed set of alternatives, the triage method allows the set of alternatives to increase over time. By partitioning the set of alternatives into three groups transfer, reject, pending the decision maker can focus more involved decision efforts on alternatives that have a better chance of satisfying the overall decision objective. Further, to account for changing preferences, the triage method uses a time-dependent function that gives pending alternatives a chance to be reevaluated.

Many screening methods use comparison cutoff values without considering value trade-offs. The triage method, however, combines the benefits of comparison screening with value trade-offs by using a value hierarchy to score incoming alternatives and compare their values to hypothetical ‘needed’ and ‘desired’ alternatives. Because every value hierarchy has inherent uncertainty due to the weights, the triage method employs intervals to account for these uncertainties and disagreements. Theorem 1 provides an easy way to calculate a spectrum of values for each alternative that summarizes their respective strengths and weaknesses. In summary, the triage method can be used to reduce alternatives for further analysis which can include choosing the best portfolio or the best single decision. After screening (stage 1 of the decision process), the transferred alternative would move into a choice model (stage 2) where an in-depth analysis would take place. The literature offers numerous choice models. For example, in portfolio analysis, the decision maker could implement ideas from [57], [63], [7], or [21].

While the Triage Method accommodates situations when the decision maker can neither create the alternatives nor analyze all alternatives at once, it can be used in static

scenarios. The decision maker could specify either a ‘needed’ or ‘desired’ alternative and apply the Triage Function without time to partition the alternatives.

3.6 Theorem Proof

Theorem 2. Let $S^* = \{s_1^*, s_2^*, \dots, s_n^*\}$ with $s_i^* \in [-\lambda_i w_i, \lambda_i w_i]$ and $\sum_{i=1}^n s_i^* = 0$ be the set that maximizes $v(A)$. i.e.

$$\max_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i + s_i^*) v_i(a_i).$$

Then,

$$\min_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i - s_i^*) v_i(a_i). \quad (3.6)$$

Equivalently, $\vec{w}_{min} = (w_1 - s_1^*, w_2 - s_2^*, \dots, w_n - s_n^*)$.

Proof. Proof Suppose this is not true. Then there exists a set $T = \{t_1, t_2, \dots, t_n\}$, $T \neq S^*$, with $t_i \in [-\lambda_i w_i, \lambda_i w_i]$ and $\sum_{i=1}^n t_i = 0$, such that

$$\min_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i - t_i) v_i(a_i).$$

However, if this is true, then

$$\begin{aligned} \sum_{i=1}^n (w_i - t_i) v_i(a_i) &< \sum_{i=1}^n (w_i - s_i^*) v_i(a_i) \Rightarrow \\ - \sum_{i=1}^n t_i v_i(a_i) &< - \sum_{i=1}^n s_i^* v_i(a_i) \Rightarrow \\ \sum_{i=1}^n t_i v_i(a_i) &> \sum_{i=1}^n s_i^* v_i(a_i) \Rightarrow \\ \sum_{i=1}^n w_i v_i(a_i) + \sum_{i=1}^n t_i v_i(a_i) &> \sum_{i=1}^n w_i v_i(a_i) + \sum_{i=1}^n s_i^* v_i(a_i) \Rightarrow \\ \sum_{i=1}^n (w_i + t_i) v_i(a_i) &> \sum_{i=1}^n (w_i + s_i^*) v_i(a_i). \end{aligned}$$

This contradicts the fact that S^* maximizes $v(A)$. Therefore,

$$\min_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i - s_i^*) v_i(a_i).$$

□

IV. Continuous Decision Support

*The contents of this chapter were published in 2014 in **The International Journal of Multicriteria Decision Making**. They have been reformatted to comply with the AFIT style guide.*

This paper presents extensions to the Triage Method for addressing continuous decision problems. These provide decision makers more tools with which to address situations where alternatives present themselves over time. The original Triage Method provides a criteria with which to divide these alternatives into those that should be selected, those that should be rejected, and those that should be held pending until a later decision epoch. The proposed extensions offer different criteria for accomplishing this partition. A sample problem is introduced to compare the effectiveness of the various methods and avenues for further improvements are discussed.

4.1 Introduction

[45] classically defines a decision as “an irrevocable allocation of resources.” [79] extends this definition to “a conscious, irrevocable allocation of resources with the purpose of achieving a desired objective.” [66] adds more specificity, defining a decision as “a choice out of a number of alternatives...made in such a way that the preferred alternative is the ‘best’ among the possible candidates.” [59] eschews a formal definition altogether and instead details aspects of a decision: available alternatives, differing outcomes, and uncertainty as to which outcomes are associated with the alternatives.

As these definitions move from general to more and more specific, they begin to more closely resemble the activity that most people associate with decision-making: choosing from a set of known alternatives. Indeed, [52] state that decision analysis “is designed to help the individual make a choice among a set of prespecified alternatives.” But does this

represent the full spectrum of decisions that decision makers are called upon to make? Arguably no, it merely represents the most analytically tractable of common decision situations. The situation can be complicated by a number of factors including, though certainly not limited to:

- The decision maker may seek to choose more than one alternative
- The decision maker may have the alternative to defer a decision
- The decision maker may be held to multiple resource constraints
- The decision maker may expect other alternatives will present themselves later

It is not uncommon for a decision maker to face a situation that is complicated by all of the above factors in what we term a *continuous* decision problem. Such problems are characterized by the following conditions:

4.1.1 The Prospect of Multiple Decision Epochs.

A key difference between a continuous decision situation and a traditional one is the decision maker's reasonable expectation that the ultimate decision will be achieved via a series of multiple, smaller decisions rather than a single, monolithic decision. In a continuous decision problem, the decision maker has the option to select from the available alternatives or to defer a decision with the clear expectation that they will revisit the problem in the future. The decision maker may choose from the set of known alternatives at time t , or defer a decision until $t + x$. At time $t + x$, the decision maker again may choose from the known alternatives, or defer until $t + y$, $y > x$. The expectation is that the decision environment will be more favorable at that time. Perhaps new and better alternatives will be available, or the uncertainty surrounding the known alternatives will have decreased. Perhaps the decision maker will have a better understanding of the requirements associated with the decision problem and can thus make a more informed choice. Or perhaps more

resources will have become available, rendering previously infeasible alternatives feasible. Of note, the passage of time may preclude previously available alternatives from being pursued. Whatever the motivation, the expectation of revisiting the problem marks the distinction.

This is a subtle difference, and it may be argued that this is no different from a series of traditional “one-time” decisions and may be approached as such with traditional methods. Most traditional methods however are geared toward a definition similar to Lootsma’s. If presented with a collection of inferior alternatives, they will help choose the least inferior, but may not guide the decision maker to defer. [38] The decision may be revisited later, but if the selected alternative is changed, the resources expended to date may well have been wasted. Frequently engaging in such revisions may prove costly.

4.1.2 Availability of Alternatives and Data.

Whether the decision maker is seeking a single feasible alternative or to assemble a collection that is in some sense optimal, the difficulty is compounded if there is a reasonable expectation that the entire set of alternatives is not available, or that substantially more information about the decision problem will become available at some point in the future. Indeed, this is in all likelihood the reason the decision maker might choose to defer a decision. While the expectation may not, in the end, be realized, its presence is another defining condition, altering the decision problem from a traditional choice to a continuous decision problem.

The introduction of this element makes the distinction between analysis of traditional decision problems and continuous ones clearer. As stated earlier, traditional methods of decision analysis tend to assume that the complete set of alternatives is defined, and that the decision maker only needs assistance in identifying the best option or set of options. Continuous decision problems require analytical techniques that can provide insight to decision makers in the absence of this assumption.

Consider the classic case of constructing a portfolio of research projects. [62] In a traditional approach, the decision maker would collect potential projects until some deadline, and then evaluate them all via their preferred methodology. Resource constraints could be treated via linear programming techniques to select a subset that was optimal according to some measure. Once selected, the portfolio would be considered complete. The decision maker may then collect new proposals until a future date, at which time they could choose to evaluate the new proposals to construct an additional portfolio, or re-evaluate the entire set of known proposals. If the entire set is re-evaluated, the decision maker may face the prospect of choosing between abandoning a previously started project or accepting the “sub-optimality” of continuing.

In a continuous approach, the decision maker would evaluate each proposal as it was presented, making a series of smaller decisions about each individual project. At each decision epoch then, the decision maker is faced with not only evaluating the merits of the new alternative, but considering the possibility that a better alternative might become available later. There is still the possibility that previously selected alternatives may no longer be part of the optimal set, but since the decision maker is only considering a single new alternative at a time, the decision becomes a smaller (though not necessarily easy) tradeoff to consider.

While the resulting portfolios will likely be different, the ultimate goal of an optimal portfolio remains unchanged. Almost by definition, the portfolio constructed via the continuous framework will be “less optimal” than one constructed when the entire set of alternatives is known. This is to be expected, as decisions made in the traditional construct benefit from more complete information. The desire on the part of the decision maker, and the aim of the analyst, is to gain sufficient flexibility to offset the loss of optimality.

4.1.3 A Finite Time Horizon.

While theoretically a decision may be deferred indefinitely, for practical applications an alternative or alternatives must eventually be chosen. Further, the existence of a finite time horizon may facilitate analytical methods that provide the decision maker insight. In many cases, the resources provided to the decision maker have an “expiration date” beyond which they are no longer available to be allocated toward alternatives. As this date approaches, the decision maker may (or may not) be willing to modify their decision criteria. Continuous decision problems require analytical techniques that can reflect this facet of the decision maker’s thinking.

Taken together then, these three elements lead to the definition of a continuous decision problem as one in which the decision maker, within a finite time horizon, expects to sequentially engage in more than one decision event en route to a final selection. At each decision epoch, the decision maker has the ability to make zero or more selections from known alternatives, or to defer a decision until a later date when the set of alternatives may have changed.

4.2 Background and Related Work

Our key focus here is in selecting items to form a portfolio of some type. The vast majority of the literature on this topic proceeds from the assumption that the set of alternatives is known to the decision maker. [27] summarizes a number of approaches with the assumption that the set of alternatives is known being so fundamental as to never be explicitly acknowledged. Others such as [44] focus on identifying an appropriate set of attributes that can be identified in order to rank projects, an activity that is only meaningful if the set is known. [43] deal with the construction of the objectives hierarchy and selection of attributes, and are explicit in acknowledging the decision maker’s knowledge of the set of alternatives.

[76] suggests an approach in which two criteria types are used in a multiple criteria model: those that “are used to characterize both the entire portfolio and its individual elements,” and those that “are solely used to evaluate the portfolio as a whole but not its elements.” The first set is assumed to be composed of quantitative measures, and is given priority over the second set which may consist of more qualitative criteria. [19] focuses on methods for evaluating the portfolio as a whole rather than its component pieces, as “the combination of individually good projects unnecessarily constitutes the optimal portfolio.” This is a common refrain in portfolio selection problems, but again does not address the sequential evaluation of projects.

[64] investigate a search problem where the decision maker evaluates a sequence of options with the aim of selecting a single “best” alternative. In this construct, with each alternative the decision maker encounters they must decide whether to select the option and terminate the search, reject the option and continue the search, purchase more information about an attribute of the option, or terminate the search in favor of a status quo option. The authors provide a dynamic programming approach to this problem and offer a method for determining optimal policies. The potential for adapting this approach to the construction of a resource constrained portfolio of options is a potential avenue for further research.

Perhaps the most expansive body of related research in the literature deals with constructing portfolios that maximize financial gain. These may be pharmaceutical R&D projects as in [84], stocks and other securities as in [61], or oil field explorations as in [13]. None of these approaches capture all the aspects of the continuous decision problem as described.

In most R&D project evaluation methods the sequential decisions are not the acceptance/rejection of different projects, but rather whether or not to continue a project based on the information gained in the current stage of research. This relates them to the sequential search problem in the sense that the decision is whether to continue to purchase

additional information about the option or terminate consideration of that option. Bayesian approaches are particularly attractive in this setting. Exploration problems tend to assume again that the set of options is known and focus on methods for reducing the exponential growth of pairwise comparisons necessary to form a joint probability distribution. This distribution is then used to order (i.e. rank) the exploration opportunities. The application of financial analysis methods, particularly real options methods, is a growing area of interest, though again it usually assumes that the entire set of alternatives is known. [35, 80]

The most directly applicable approach identified to date is the Triage Method described by [38]. The Triage method uses global sensitivity analysis to compare a single alternative's best and worst case performance potential to a cutoff value α . The method begins with the specification of a linear additive hierarchical value model as would be used in a classic decision context. Using whatever means the analyst chooses to elicit information from the decision maker, a value hierarchy is constructed. Weights for each element of the value hierarchy are specified, and single dimensional value functions for each element are determined. Finally, sensitivity intervals are specified for each of the weights in the value hierarchy. This allows the decision maker to express any uncertainty they may feel toward the hierarchy's weighting. For example, element 1 in the hierarchy may be assigned a global weight of 0.2, but the decision maker may specify that the weight could actually be anywhere between 0.15 and 0.25.

To this point, the model has been constructed no differently than it would be in a classic decision case where all alternatives were available for a single decision activity. The first difference is encountered when the analyst assists the decision maker in specifying a cutoff value α . The second, and key, difference is the temporally distributed arrival of alternatives. Now, as an alternative A arrives, its performance on each of the single dimension value functions is assessed, weights are applied, and an initial value score $V(A)$ is determined. Linear programming methods are then used to determine the set of hierarchy weights W_{max}

that maximize $V(A)$ while staying within both the sensitivity intervals specified by the decision maker and the overall constraint that the sum of the hierarchy weights must equal 1.

At this point we have determined the maximum value $V_{max}(A)$ that alternative A can achieve in the given model construct. Gutman and Weir offer a proof in their paper that given $V_{max}(A) = V(A) + \Delta(A)$ then $V_{min}(A) = V(A) - \Delta(A)$. There is no need to re-run the linear program in order to minimize $V(A)$, $V_{min}(A)$ can be determined simultaneously with no appreciable computational cost. Gutman and Weir's proof of this reciprocal calculation is reproduced in Appendix 4.6.

The Triage method now compares the calculated scores $V_{max}(A)$ and $V_{min}(A)$ to the cutoff value α and the alternative is triaged into one of three categories:

1. Where $V_{min}(A) > \alpha$ the alternative is selected
2. Where $V_{max}(A) \leq \alpha$ the alternative is rejected
3. Otherwise, the alternative is held pending further analysis

The Triage Method provides a fast, analytically defensible method for continuously evaluating alternatives as they arrive. Unlike methods applied to static decision situations, where the ranking of an alternative's value score relative to other alternatives is the focus, this method keys on the value of each individual alternative. Each alternative is allowed to reach its maximum potential value score within the weight space specified by the decision maker. Since the set of weights that maximize one alternative's value are likely different from those that maximize another's, it is not useful to compare the maximum value scores to one another. Instead, they are compared to the cutoff value α . [38] also propose a value of α that changes over time. This modification will be discussed in Section 4.5. A further open question is the best method for establishing the value of α .

As initially proposed, the Triage Method is primarily used as a screening tool to rapidly evaluate a stream of alternatives and determine which are promising candidates for further, more detailed analysis. This is highly applicable in a case where the final decision is the selection of a portfolio of alternatives as opposed to a single alternative. An obvious weakness in applying the methodology to such problems is that the Triage method does not directly address resource constraints. A desirable extension to the methodology would be to provide for consideration of constraints and thus allow its use for portfolio selection in addition to screening.

4.3 Methodological Extensions

We examine two possible extensions to the Triage Method that provide greater potential for use in selection, as opposed to screening. In Section 4.4 these will be illustrated by way of a simple decision model.

The original Triage method looks only at the maximum and minimum value scores an alternative can achieve relative to the cutoff value α . It is often the case in a portfolio selection problem that the decision is limited by one or more constraints. [57] For instance, the total cost of the selected options cannot exceed a stated budget, or the number of man-hours available in each of a variety of technical specialties is limited. There is a significant difference between two alternatives whose V_{min} scores exceed α if one consumes a substantial portion of available resources while the other's resource requirements are more modest. As originally constructed the Triage method does not make this distinction.

The first extension, which is termed Triage+, introduces the consideration of constraints into the Triage methodology. With this extension, rather than simply considering an alternative's maximum and minimum value scores, the ratio of those scores to a constraint is considered. In the case where there is a single constraint the ratio is apparent. In the case of multiple constraints, further analysis may be necessary to determine whether there is a single constraint that is most likely to be binding. In this method, the

cutoff value α_+ now represents a benefit/constraint ratio instead of an absolute value score. As such, it is no longer bound to values in the interval $[0, 1]$ and the selection of its value may be more difficult than with the standard α . In Section 4.5 we will discuss the difficulties in both identifying an appropriate constraint ratio and setting an α_+ value.

The Triage+ method begins as the Triage method does, with the construction of a value model, single dimension value functions, weights, and sensitivity intervals. In addition, one or more constraints C that will restrict the construction of the portfolio are identified and quantified. One of these constraints C_{crit} is identified as the most critical and will be used to form a ratio that is compared to α_+ . Finally, a value for α_+ is defined. As each alternative A arrives, the method proceeds as follows:

1. Compare the alternative's resource requirements to the current level of available resources. If selecting the alternative would violate resource constraints, reject the alternative. If the alternative is feasible, continue the evaluation.
2. Score alternative A against the value model and determine $V(A)$
3. Determine $V_{max}(A)$ and $V_{min}(A)$ as in the Triage method
4. Form the ratios $R_{max}(A) = \frac{V_{max}(A)}{C_{crit}}$ and $R_{min}(A) = \frac{V_{min}(A)}{C_{crit}}$
5. Compare these values to α_+
 - (a) Where $R_{min}(A) > \alpha_+$ the alternative is selected
 - (b) Where $R_{max}(A) \leq \alpha_+$ the alternative is rejected
 - (c) Otherwise, the alternative is placed in a pending category
6. If the alternative was selected, decrement the available resource levels according to the alternative's requirements
7. The final portfolio is constructed by:

- (a) Use binary integer programming to select the set of pending alternatives with the maximum sum of value scores, subject to the remaining resource constraints
- (b) Combine the selections from the pending alternatives with those that were selected outright

The net effect of this extension is to reduce the likelihood that the model will select alternatives which are “resource-intensive” particularly with regard to the scarcest resources. This is consistent with the decision process in most portfolio selection cases as a decision maker would be less likely to select an alternative which scores well but consumes a large portion of the available resources, preferring instead those that offer a greater return. As highlighted above, this is a significant departure from the original Triage formulation which, since it was only interested in screening, would select the first alternative with $V_{min} > \alpha$ even if its resource requirements consumed the entire available budget and resulted in a portfolio of one alternative.

The second proposed extension, termed Triage++, builds upon the first by introducing a temporal element to the Triage+ considerations. This requires either more familiarity with the decision situation or the establishment of assumptions about the distribution of alternative scores and their arrivals. An approximation of the CDF cumulative distribution function (CDF) for the Triage+ ratios described above is developed, as is an estimate for the average number of alternatives arriving per time period. When a new alternative arrives, it is scored as in the Triage+ method. The CDF for these ratios is then consulted to determine what percentage of future alternatives should be expected to exceed the ratio of the current alternative. This percentage is multiplied by the number of remaining time periods and the average arrival rate per period to arrive at an expected number of better alternatives before the end of the decision cycle. In this method, the cutoff value α_{++} now represents the number of alternatives the decision maker can expect to see before the end of the time horizon whose benefit/constraint ratio exceeds the current alternative’s.

For example, let us assume that over the next month we will be evaluating proposals for research projects to fund. This is an activity that our organization has engaged in numerous times in the past. Based on historical data, we make the assumption that we will receive an average of λ proposals per day. Further, we have used our historical data on the value scores of proposals to derive an empirical CDF for value scores and the ratio of value score to funding required. Thus our critical constraint C_{crit} is the funding required by a proposal. As each alternative A arrives, the method proceeds as follows:

1. Compare the alternative's resource requirements to the current level of available resources. If selecting the alternative would violate resource constraints, reject the alternative. If the alternative is feasible, continue the evaluation.
2. Score alternative A against the value model and determine $V(A)$
3. Determine $V_{max}(A)$ and $V_{min}(A)$ as in the Triage method
4. Form the ratios $R_{max}(A) = \frac{V_{max}(A)}{C_{crit}}$ and $R_{min}(A) = \frac{V_{min}(A)}{C_{crit}}$ as in the Triage+ method
5. Consult the CDF to determine $P(R_{max}(A))$ and $P(R_{min}(A))$
6. The decision maker would expect to see N_{max} alternatives with ratios better than $R_{min}(A)$ in the remaining T time periods (days in this case)
 - (a) $N_{max} = (1 - P(R_{min}(A))) \cdot \lambda \cdot T$
7. The decision maker would expect to see N_{min} alternatives with ratios better than $R_{max}(A)$ in the remaining T time periods (days in this case)
 - (a) $N_{min} = (1 - P(R_{max}(A))) \cdot \lambda \cdot T$
8. Compare these values to α_{++}
 - (a) Where $N_{max} \leq \alpha_{++}$ the alternative is selected

- (b) Where $N_{min} > \alpha_{++}$ the alternative is rejected
 - (c) Otherwise, the alternative is placed in a pending category
9. If the alternative was selected, decrement the available resource levels according to the alternative's requirements
10. The final portfolio is constructed by:
- (a) Use binary integer programming to select the set of pending alternatives with the maximum sum of value scores, subject to the remaining resource constraints
 - (b) Combine the selections from the pending alternatives with those that were selected outright

The calculations and comparison made to α_{++} may not be intuitive. In short, this process attempts to quantify the decision maker's expectation of seeing alternatives that are better than the one currently under consideration before the conclusion of the decision cycle. In this case "better" is defined as having a greater value score to critical resource ratio, as in the Triage+ extension. Given that the decision maker is currently considering an alternative whose value score to critical resource ratio is in the range $[R_{min}, R_{max}]$, they could expect to see more alternatives in the future whose ratios exceed R_{min} than whose ratios exceed R_{max} . Thus the greatest number of "better" alternatives they can expect over the next T days is $N_{max} = (1 - P(R_{min}(A))) \cdot \lambda \cdot T$ and the smallest number of "better" alternatives they can expect is $N_{min} = (1 - P(R_{max}(A))) \cdot \lambda \cdot T$. If $N_{min} > \alpha_{++}$ then the decision maker should expect see at least α_{++} better alternatives before the close of the decision cycle and so the current alternative can be passed over to wait for a better one. On the other hand, if $N_{max} \leq \alpha_{++}$ then the greatest number of better alternatives that can be expected does not exceed α_{++} and so the current alternative is selected.

4.4 An Illustrative Example

The Triage Method itself, and particularly the proposed extensions, are best illustrated with a sample application. The problem and value model shown here are contrived for simplicity and real-world applications are likely to be more complex, but the fundamental approach remains unchanged.

We consider the case of a sports memorabilia shop whose owner is an avid personal collector of baseball cards. Customers come to his store to buy and sell classic baseball cards. The owner is also building a baseball card display to enter in a collector's competition 15 days from now. Between now and then, as customers bring in antique baseball cards, he must decide which ones to purchase for his personal display, and which to add to the store's inventory of cards for sale.

Classic baseball cards can vary in size. The most common size is 8.75 in², but a small number of early baseball cards are either 5.25 in², or 9.97 in². The competition rules limit the size of the display to a total of 100 square inches of card. The baseball card market has a well known pricing structure, and "book values" for mint condition cards are published. The actual price paid when buying and selling a card is determined by the parties involved, but is based on the book value and the condition of the individual card. The owner will enter a competitive category where the total book value of cards in the display cannot exceed \$5,000.

Based on his experience, the owner has developed a model for how he believes judges will evaluate cards, with the total "quality" of the display being the sum total of the evaluation of each card included. The model is detailed below. The dominant aspect is the book value of the card, which is the dollar value listed in an authoritative pricing guide such as Beckett's for a mint condition example of the given card. These values serve as the pricing guide for the market and take into account the relative rarity of the card, the quality

of the player, etc. By using this criteria the owner can avoid attempting to model all these various facets over the large body of baseball cards on the market.

Price is not the only factor that influences a judge's opinion however. Older cards are always impressive and can catch a judge's eye even if, perhaps due to there being a greater supply of the card, its book value is lower than the age might suggest. Finally, the condition of the card is important. The book value assumes a card in mint condition, and a card in such condition will receive full points in this criterion. As the condition degrades, so too do the points awarded in this criterion.

The challenge is to develop a methodology that will assist the owner in determining, as customers bring in cards, which ones he should acquire for his personal display and which he should place in the store's for-sale inventory.




Figure 4.1: Baseball Card Value Hierarchy

In order to explore the limitations of current methodologies and highlight the potential impacts of new techniques, a simulation was developed to model the problem as described above, with a few simplifying assumptions made for time considerations:

- The number of cards brought in for sale to the store each day is a random variable that is Poisson distributed with $\lambda_{in} = 2$.




Figure 4.2: Value Curve for Baseball Card Book Value




Figure 4.3: Value Curve for Baseball Card Age

- The number of cards purchased from the shop each day is a random variable that is Poisson distributed with $\lambda_{out} = 2$.
- 20% of cards are 5.25 in^2 , 10% are 9.97 in^2 , and 70% are 8.75 in^2




Figure 4.4: Value Curve for Baseball Card Condition

- The age of a card is drawn as an exponential random variable with $\lambda_{age} = 40$ years. Values less than 1 are forced to be 1, and values greater than 100 are forced to be 100.
- The book value of a card is an exponential random variable with λ_{value} assigned in the range [\$500, \$2750] depending on the age of the card. Older cards have higher average book values.
- A card's condition is determined by drawing a uniform [0, 1] random number and comparing it to threshold values based on the card's age. "Poor" condition cards have their sale price set at 10% of book value $\pm \$100$, "Fair" condition cards have theirs set at 25% of book value $\pm \$100$, etc. Older cards are more likely to be in poorer condition.
- The shop has an initial inventory of 5 cards for sale.
- No cards that are brought in are turned away; the only question is whether or not they are placed in the for-sale inventory or the display collection. The purchase price of a

card is its sale price multiplied by a discount factor, uniformly drawn from the range [0.75, 0.95].

- The weights in the value hierarchy may vary by as much as $\pm 25\%$ for sensitivity analysis.

We explore the problem via five potential methodologies:

- Buy All-Sell None: This is an unrealistic option that provides an upper bound as to the value the display can achieve. The owner purchases every card that is brought into the store for use in his personal display and holds them until the day before the competition, at which time he assembles the highest scoring display possible using the original weights in his value hierarchy. No card sales take place in this approach.
- Buy and Sell: In this approach, the owner assembles the best display possible with the cards on hand at the end of each day. The following day, the cards that were selected for the display are held out of the for-sale inventory.
- Triage: Using the Triage method the owner evaluates each card as it comes in. The day before the competition, he assembles the best possible display from the cards that have been selected for the display and the cards that were held as pending.
- Triage+: Identical to the Triage approach, only instead of comparing the card's value score to α , the card's $\frac{\text{value score}}{\text{book value}} \times 1000$ ratio is compared to α_+ .
- Triage++: Based on the assumptions that were made for the distribution of card attributes, 10,000 cards were generated. The cumulative density (CDF) of the $\frac{\text{value score}}{\text{book value}} \times 1000$ ratio was determined by using MatLab to fit a curve to the resulting data. As each card comes in, its ratio is calculated. Then, based on the CDF and the number of days remaining before the competition, an estimate is formed as to how

many cards can be expected to come in that are better than the current card, and this is compared to α_{++} .

A MatLab simulation was constructed to execute 100 iterations at a variety of values for α , α_+ , and α_{++} . In each iteration, an identical stream of incoming cards and card sale events was presented to each methodology. Because each methodology may make different decisions about which cards may be offered for sale, each method maintains its own, independent inventory. When a sale event occurs, each method chooses a card for sale from its own inventory. The choice is random, but is weighted to make less expensive cards more likely to sell.

4.4.1 Triage Results.

The Triage Method was applied with values of α ranging from 0.25 to 1 in increments of 0.0125. The average portfolio values achieved for each α value are depicted in Figure 4.5, as is the average value achieved by the Buy All-Sell None and Buy and Sell approaches. As alluded to earlier, the Buy All approach provides an upper bound to the obtainable portfolio value, and in this case averages 4.48, while the Buy Sell approach achieves an average portfolio value of 3.57. The Buy Sell approach represents a more realistic approach as the average cost (loss) of pursuing the Buy All approach was \$22,375 with no offsetting income from card sales. The Buy Sell approach, on average, had a net cost of \$363.

The Triage Method performed best with values of α in the range [0.275, 0.425]. Within this range, the average portfolio value exceeded that achieved by the Buy Sell method. Outside these bounds performance diminished rapidly. Values of α in the range [0.325, 0.3625] offered performance near that of the Buy All approach. This performance comes at a significant financial cost however, as can be seen in Figure 4.6, with the card shop experiencing a net cost in excess of \$5,000. What is essentially happening is that a great many of the cards are being triaged into the “pending” category and thus are not available




Figure 4.5: Triage Portfolio Value

for sale when a card sale event is encountered. In the final portfolio construction these cards are available for consideration, thus leading to the improved portfolio value scores, but at the cost of not having them available for sale. In fact, the Triage method provides a financial loss for almost all values of α , save only those at the extreme tail, where portfolio value performance is abysmal.

4.4.2 Triage+ Results.

The Triage+ method provided performance similar in some ways to the Triage method, shown in Figure 4.7. Like α , the most effective values of α_+ lay in about 25% of the evaluated range near the lower bound. In this case that represents values in the range [0.25, 0.85]. Within this range average portfolio value exceeded that of the Buy Sell method, and beyond this range performance again dropped off sharply. The decrease in performance was not as steep as the Triage method's however.

Where the Triage+ method outperformed Triage is in the financial picture behind the portfolio, as shown in Figure 4.8. For α_+ values of 0.45 and higher, the Triage+ method




Figure 4.6: Triage Net Profit/Loss




Figure 4.7: Triage+ Portfolio Value

delivered a net profit to the card shop. This stands in stark contrast to the Triage method which did not deliver a profitable performance until α values were near 1, at which point portfolio values were well below 1. Overlaying these two performance measures, we find an interval of approximately $[0.45, 0.85]$ for α_+ where the method was both profitable and achieved portfolio values that exceeded the Buy Sell approach. Triage+ achieves this performance by holding fewer cards in the Pending category, and thus providing greater opportunity to sell valuable cards when sales events occur.




Figure 4.8: Triage+ Net Profit/Loss

4.4.3 Triage++ Results.

Triage++ did not match the peak performance of Triage+, but it did provide performance superior to the Buy Sell method over a much larger range of α values as depicted in Figure 4.9. This is important because, as described earlier, how to select the best value of α is an open question. For α_{++} values greater than 3.4 the method consistently delivered average portfolio values greater than that achieved by the Buy Sell method.

Knowing that portfolio value performance is relatively insensitive to the choice of α_{++} value simplifies the selection of a value by allowing the analyst to focus on the financial performance when choosing a value.




Figure 4.9: Triage++ Portfolio Value

Triage++ is not as stable as Triage+ in terms of financial performance as can be seen in Figure 4.10. Net profit peaks when α_{++} is 3.6 and begins to steadily decline after that, falling below the performance of the Buy Sell method for most values from 8.2 onward. There is also considerably more variability in the financial performance of the Triage++ method than the Triage+ method. Triage++ does still present a larger range of α values where both financial and portfolio quality measures exceed those of the Buy Sell method however. Figures 4.11 and 4.12 below combine the portfolio value and financial performance charts for the Triage+ and Triage++ methods respectively. In each figure, the shaded region represents the range of α values where the method outperforms the Buy Sell method in both portfolio value and net profit/loss.




Figure 4.10: Triage++ Net Profit/Loss




Figure 4.11: Triage+ Combined Performance




Figure 4.12: Triage++ Combined Performance

4.5 Conclusions and Future Work

The purpose of this paper is to highlight the Triage Method and to offer two possible extensions for use in addressing continuous decision problems. The Triage+ method is the more easily applied, while the Triage++ method requires either greater familiarity with the decision problem or the development of additional assumptions. These extensions both lend themselves to further research questions, as does the prospect of further extensions to the basic Triage methodology.

An obvious first question is what aid the decision analyst can provide the decision maker in specifying the value of α for the different methodologies. The performance of the various techniques varied greatly with the α value applied even in our simple example and there is no reason to believe this sensitivity is reduced in more complex decision environments. An ideal approach would be to develop some function of the resource

requirements of the alternatives to calculate an α value for each alternative based on its requirements relative to the constraints.

This in turn leads to two further considerations. First, as alluded to earlier, it may be desirable to vary the value of α with time. As the end of the decision epoch approaches for example, the decision maker may wish to relax the α value to admit previously excluded options rather than risk reaching the end of the epoch with unallocated resources. Second, the performance of the extensions under more complex constraint scenarios must be investigated. In the example presented here, the constraint picture was simple by design, and an appropriate ratio for Triage+ and Triage++ was apparent by inspection. This is likely not the case in a more realistic, complex decision environment.

Finally, the Triage++ extension is dependent on the specification of a CDF for the ratios under consideration. In a real world problem this is almost certain to require the application of a number of assumptions, and overall performance will be heavily dependent on their quality. It is worth investigating both methods for determining these, as well as the prospect that some assumed distribution may be more robust than others. If such distributions can be identified they will prove quite useful in situations where historical data regarding the problem at hand is scarce.

4.6 Theorem Proof

Theorem 3. *Let $S^* = \{s_1^*, s_2^*, \dots, s_n^*\}$ with $s_i^* \in [-\lambda_i w_i, \lambda_i w_i]$ and $\sum_{i=1}^n s_i^* = 0$ be the set that maximizes $v(A)$. i.e.*

$$\max_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i + s_i^*) v_i(a_i).$$

Then,

$$\min_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i - s_i^*) v_i(a_i). \quad (4.1)$$

Equivalently, $\vec{w}_{min} = (w_1 - s_1^, w_2 - s_2^*, \dots, w_n - s_n^*)$.*

Proof. Proof Suppose this is not true. Then there exists a set $T = \{t_1, t_2, \dots, t_n\}$, $T \neq S^*$, with $t_i \in [-\lambda_i w_i, \lambda_i w_i]$ and $\sum_{i=1}^n t_i = 0$, such that

$$\min_{\vec{w} \in W_S} (v(A)) = \sum_{i=1}^n (w_i - t_i) v_i(a_i).$$

However, if this is true, then

$$\begin{aligned} \sum_{i=1}^n (w_i - t_i) v_i(a_i) &< \sum_{i=1}^n (w_i - s_i^*) v_i(a_i) \Rightarrow \\ &-\sum_{i=1}^n t_i v_i(a_i) < -\sum_{i=1}^n s_i^* v_i(a_i) \Rightarrow \\ &\sum_{i=1}^n t_i v_i(a_i) > \sum_{i=1}^n s_i^* v_i(a_i) \Rightarrow \\ \sum_{i=1}^n w_i v_i(a_i) + \sum_{i=1}^n t_i v_i(a_i) &> \sum_{i=1}^n w_i v_i(a_i) + \sum_{i=1}^n s_i^* v_i(a_i) \Rightarrow \\ \sum_{i=1}^n (w_i + t_i) v_i(a_i) &> \sum_{i=1}^n (w_i + s_i^*) v_i(a_i). \end{aligned}$$

This contradicts the fact that S^* maximizes $v(A)$. Therefore,

$$\min_{w \in W_S} (v(A)) = \sum_{i=1}^n (w_i - s_i^*) v_i(a_i).$$

□

V. The Effects of Decision Problem Parameters on the Effectiveness of the Extended Triage Method

*The contents of this chapter were submitted to **The International Journal of Multicriteria Decision Making** in September 2015. They have been reformatted to comply with the AFIT style guide.*

We explore the effectiveness of the extended Triage Method when applied to a realistic data set representing a real-world decision situation. We analyze the impact of various factors of the decision environment on the effectiveness of the Triage extensions as well as the trade-off between temporal flexibility and portfolio quality. We begin to identify guidelines to assist decision makers in employing the Triage extensions.

5.1 Introduction

In an earlier work we introduced the Triage Method as a way to use multi-objective decision analysis techniques to screen decision alternatives that arrived asynchronously over time. [86] As originally formulated, the Triage Method operated without regard to resource constraints to partition the set of alternatives into one of three sets: those that are accepted for further consideration, those that are rejected outright, and those that are held pending for potential consideration in the future. As such Triage functioned as a screening process that could form the basis of a multi-stage selection process such as described in [5] but was not suited to the task of making resource-constrained selections. We then offered extensions to the Triage Method that allowed it to consider resource constraints in selecting portfolio members from a set of alternatives that arrived over time, and demonstrated the potential of these extensions on a small sample problem. [41]

We turn now to further investigating the effectiveness of the Triage extensions when applied to a more complex, realistic decision environment. Further, we seek to examine

which elements of the decision environment most affect the extensions so that we may provide decision makers guidance on how to effectively employ them. In doing so, we investigate measures of the tradeoff between temporal flexibility and final portfolio value.

Before examining these topics, we first define the decision environment for which the Triage extensions are suited, one which we term a “continuous decision problem”. As described in [41], a continuous decision problem is characterized as one in which the decision maker, within a finite time horizon, expects to assemble a portfolio by sequentially engaging in more than one decision event en route to a final selection. At each decision epoch, the decision maker has the ability to make zero or more selections from known alternatives, or to defer a decision until a later date when the set of alternatives may have changed.

Consider the classic case of constructing a portfolio of research projects. [62] In a traditional approach, the decision maker would collect potential projects until some deadline, and then evaluate them all via their preferred methodology. In [75] the authors suggest these are typically either a corporate finance perspective such as net present value, an operations research perspective that treats the issue as a knapsack problem, or a decision analysis perspective that applies decision trees or multi-criteria methods to rank alternatives. In any case, there is an implicit assumption that the entire set of alternatives is known at decision time. [72] [2] Indeed, in [56] the authors go so far as to explicitly state, “We will assume that the set of alternatives $A_j, j = 1, \dots, J$ has been identified for the decision problem,” while in [18] they state “the task of multiple criteria decision analysis is to help a decision maker choose, rank or sort alternatives *within a finite set* according to two or more criteria.” (emphasis added). Once selected, the portfolio would be considered complete. The decision maker may then collect new proposals until a future date, at which time they could choose to evaluate the new proposals to construct an additional portfolio, or re-evaluate the entire set of known proposals. If the entire set is re-evaluated, the decision

maker may face the prospect of choosing between abandoning a previously started project or accepting the “sub-optimality” of continuing.

In a continuous approach, the decision maker would evaluate each proposal as it was presented, making a series of smaller decisions about each individual project. At each decision epoch then, the decision maker is faced with not only evaluating the merits of the new alternative, but considering the possibility that a better alternative might become available later. There is still the possibility that previously selected alternatives may no longer be part of the optimal set, but since the decision maker is only considering a single new alternative at a time, the decision becomes a smaller (though not necessarily easy) tradeoff to consider. While the resulting portfolios will likely be different, the ultimate goal of an optimal portfolio remains unchanged.

Almost by definition, the portfolio constructed via the continuous framework will be “less optimal” than one constructed when the entire set of alternatives is known. This is to be expected, as decisions made in the traditional construct benefit from more complete information. As the authors in [36] point out, in most decision problems we are not seeking a strictly *optimal* solution so much as a *satisfactory* one. The desire on the part of the decision maker, and the aim of the analyst, is to gain sufficient flexibility to offset the loss of optimality. At this stage, we assume that selected projects are fully funded and executed to completion. In [46] the authors make the point that additional flexibility may be gained by the partial funding of proposed projects, though this adds the further risk of correctly predicting the value of a partially-funded effort.

The remainder of this document is structured as follows: Section 5.2 provides a brief introduction of the Triage Method and its extensions, an overview of the historical dataset utilized in this paper, and the resulting experimental design. Section 5.3 outlines the analysis of our designed experiment to identify significant factors. Section 5.4 turns to an analysis of the overall effectiveness of the extended Triage method. Section 5.5

provides some elementary guidelines to be used in the employment of the Triage methods, and finally Section 5.6 proposes avenues for further research.

5.2 Background

5.2.1 *The Triage Method.*

We offer a brief description of the basic Triage method. For a more complete discussion see [41]. The method employs a linear additive value model and global sensitivity analysis to compare a single alternative's best and worst case performance potential to a cutoff value α . As α represents a score from the value model, it is bound to the range $[0,1]$. Using sensitivity intervals for each of the weights in the value hierarchy, a linear program (LP) is used to determine a set of weights consistent with the sensitivity intervals that maximizes the alternative's score. The nature of the additive value model ensures that the variation in score achieved by altering the weights is symmetrical. In [86] we provide a proof that if a given alternative A scores $V(A)$ in the originally specified model, and the LP-derived weights allow it to achieve a maximum score $V_{max}(A) = V(A) + \Delta(A)$, then without further processing we can determine that $V_{min}(A)$, the worst the alternative can perform within the given sensitivity intervals, is $V(A) - \Delta(A)$.

The Triage Method now compares the calculated scores $V_{max}(A)$ and $V_{min}(A)$ to the cutoff value $\alpha \in [0, 1]$ and the alternative is triaged into one of three categories:

1. Where $V_{min}(A) > \alpha$ the alternative is selected
2. Where $V_{max}(A) \leq \alpha$ the alternative is rejected
3. Otherwise, the alternative is held pending further analysis

The concept of triage is not new, and methodologies to accomplish it have been proposed. Spradlin and Kutoloski [82] propose a construct where again alternatives are separated into one of three categories: the doomed projects, the equivocal projects, and the favorite projects. The similarity to our proposed selected/rejected/pending partition is

obvious. Here again though, the implicit assumption is that the complete set of alternatives is available when the portfolio selection decision is being made.

Unlike methods applied to static decision situations, where the ranking of an alternative's value score relative to other alternatives is the focus, this method keys on the value of each individual alternative. Each alternative is allowed to reach its maximum potential value score within the weight space specified by the decision maker. Since the set of weights that maximize one alternative's value are likely different from those that maximize another's, it is not useful to compare the maximum value scores to one another. Instead, they are compared to the cutoff value α . The identification of robust methods for selecting the value of α as well as the effectiveness of allowing α to vary over time are open research questions.

5.2.2 The Triage Extensions.

Again a more complete discussion of the Triage extensions is available in [41]. We offer two fundamental approaches for extending the Triage Method to treat resource constraints and thereby become a selection method vice a strictly screening method. The first involves taking the ratio of the scores $V_{max}(A)$ and $V_{min}(A)$ to a critical resource cost, C_{crit} , of the alternative. We term this method Triage+. In the case where there is a single resource constraint the construction of this ratio is straightforward. When there are multiple constrained resources the problem becomes more complex. Potential approaches include identifying a single resource constraint that is considered most likely to be binding, or developing a method to combine resource requirements into a single measure. This ratio is now compared to a cutoff value α_+ , and triaged as in the basic method. In this extension, the value of α_+ is no longer bound to the range $[0,1]$ so effectively selecting a cutoff value is potentially more difficult.

The Triage extensions are novel, and the aim of this paper is to explore the impact of various decision environment parameters on their performance. To facilitate readability,

the following summary of the Triage+ extension is reproduced from our earlier work in [41].

As each alternative A arrives, the method proceeds as follows:

- 1. Compare the alternative's resource requirements to the current level of available resources. If selecting the alternative would violate resource constraints, reject the alternative. If the alternative is feasible, continue the evaluation.*
- 2. Score alternative A against the value model and determine $V(A)$*
- 3. Determine $V_{max}(A)$ and $V_{min}(A)$ as in the Triage method*
- 4. Form the ratios $R_{max}(A) = \frac{V_{max}(A)}{C_{crit}}$ and $R_{min}(A) = \frac{V_{min}(A)}{C_{crit}}$*
- 5. Compare these values to α_+*
 - (a) Where $R_{min}(A) > \alpha_+$ the alternative is selected*
 - (b) Where $R_{max}(A) \leq \alpha_+$ the alternative is rejected*
 - (c) Otherwise, the alternative is placed in a pending category*
- 6. If the alternative was selected, decrement the available resource levels according to the alternative's requirements*
- 7. The final portfolio is constructed by:*
 - (a) Use binary integer programming to select the set of pending alternatives with the maximum sum of value scores, subject to the remaining resource constraints*
 - (b) Combine the selections from the pending alternatives with those that were selected outright*

This extension effectively generates a benefit/cost ratio for each alternative and then compares its value to a new cutoff parameter α_+ . Alternatives which require an inordinate

amount of resources in order to achieve their level of benefit, i.e., those with a high cost relative to their benefit, are less likely to be selected than if the decision were made on benefit alone, as is the case in the original Triage Method.

The second extension, which we term Triage++, adds a temporal element to the considerations of the Triage+ method. After constructing the value model, the decision analyst also elicits estimations of the distribution of the Triage+ ratios and the rate of arrival of new alternatives, λ . A cumulative distribution function (CDF) for the Triage+ ratios is then generated. As each alternative arrives, these elements are then combined to determine an estimate for how many alternatives the decision maker can expect to arrive in the remaining time whose benefit/constraint ratio exceeds that of the current alternative.

Again, to facilitate readability, the following summary of the Triage++ extension is reproduced from our earlier work in [41].

As each alternative A arrives, the method proceeds as follows:

- 1. Compare the alternative's resource requirements to the current level of available resources. If selecting the alternative would violate resource constraints, reject the alternative. If the alternative is feasible, continue the evaluation.*
- 2. Score alternative A against the value model and determine $V(A)$*
- 3. Determine $V_{max}(A)$ and $V_{min}(A)$ as in the Triage method*
- 4. Form the ratios $R_{max}(A) = \frac{V_{max}(A)}{C_{crit}}$ and $R_{min}(A) = \frac{V_{min}(A)}{C_{crit}}$ as in the Triage+ method*
- 5. Consult the CDF to determine $P(R_{max}(A))$ and $P(R_{min}(A))$*
- 6. The decision maker would expect to see N_{max} alternatives with ratios better than $R_{min}(A)$ in the remaining T time periods*

$$(a) N_{max} = (1 - P(R_{min}(A))) \cdot \lambda \cdot T$$

7. *The decision maker would expect to see N_{min} alternatives with ratios better than $R_{max}(A)$ in the remaining T time periods*

$$(a) N_{min} = (1 - P(R_{max}(A))) \cdot \lambda \cdot T$$

8. *Compare these values to α_{++}*

(a) *Where $N_{max} \leq \alpha_{++}$ the alternative is selected*

(b) *Where $N_{min} > \alpha_{++}$ the alternative is rejected*

(c) *Otherwise, the alternative is placed in a pending category*

9. *If the alternative was selected, decrement the available resource levels according to the alternative's requirements*

10. *The final portfolio is constructed by:*

(a) *Use binary integer programming to select the set of pending alternatives with the maximum sum of value scores, subject to the remaining resource constraints*

(b) *Combine the selections from the pending alternatives with those that were selected outright*

In our original introduction of these extensions, we applied them to a simple problem with a small value hierarchy and only one binding resource constraint. We seek to further investigate the performance of the extensions using a more complex, realistic model and decision situation. We have access to five year's of data from the U.S. Air Force's Development Planning (DP) activity, and we turn to it now to provide the desired decision situation.

5.2.3 The Development Planning Model.

From FY Fiscal Year (FY) 2011 to FY 2015, Headquarters Air Force Material Command (AFMC) utilized a model-based process to prioritize DP project proposals.

Proposals were gathered over the course of a year, and at the end of the FY a portfolio of projects was selected by scoring the proposals against a value model and then using LP to select the subset that provided the greatest total value score within funding and manpower constraints. A professional military judgment (PMJ) phase was then undertaken. In the PMJ phase, alterations could be made to the final selection list to account for considerations not explicitly captured by the model. Additionally, project proposals could be “re-scoped” in order to free up resources and expand the list of chosen projects.

We do not seek to reproduce the AFMC DP experience. To begin with, we have no way of adequately modeling the PMJ phase of the selection process. Additionally, the model was improved each year based on the previous year’s experience, so we do not have a stable model over the time-frame. Finally, though DP proposals were accepted throughout each FY, they were primarily provided in response to a single data call near the end of the year. As such, we do not have data on the true temporal distribution of the project proposals. The data do however provide a very complex and realistic set on which to base our analysis.

We began by identifying nine elements of the value hierarchy that were largely common throughout the entire time-frame. In consultation with the office originally charged with managing the DP portfolio, we then rescaled the weights of the nine elements to reflect realistic values. Next we obtained the original project data for 92 unique DP proposals. This data included the proposal’s scores against the nine hierarchy elements, their direct dollar costs, and their manpower requirements in FTE full-time equivalent (FTE) positions in each of 18 different technical specialties. Finally, any missing data items were defined to reasonable values.

This formed the basis for a stochastic simulation in which project proposals could be randomly drawn (with replacement) from a pool of 92 potential projects. Portfolios could then be constructed via a variety of different techniques and the resulting values compared. Critical variables such as the arrival rate of the projects, the potential for projects to depart

while being held pending, the size of the value model, as well as the sensitivity interval and α values required by the Triage methods could all be systematically varied. We recognize that there is a limited capacity to generalize simulation results from a specific decision problem to the overall viability of any one methodology. [28] We hope to gain valuable insights though that can inform further exploration and possible applications.

5.2.3.1 Resulting Methods.

In total, we used 16 different methods to construct portfolios within the simulation:

1. **Hold All:** All arriving projects are held pending with no potential for departure. At the conclusion of the simulated time period, the optimal portfolio is constructed. This provides the upper bound for the value that could be achieved from the stream of proposals drawn.
2. **Come and Go:** Similar to the Hold All method, except there is a probabilistic chance that proposals may be withdrawn prior to the end of the simulated period. The number of withdrawals is Poisson distributed with λ_{dep} . This provides a more realistic representation of many real-world situations than the Hold All method.
3. **Random:** As each proposal arrives, there is a 50% chance it is selected and a 50% chance it is rejected.
4. **Triage:** A straight application of the Triage screening methodology with no consideration of the resource requirements of the project proposals.
5. **Triage+ Methods:** Triage+ methods as described in Section 5.2.2
 - (a) **+Cost:** Triage+ based on the ratio of value score to direct dollar cost of the proposal.
 - (b) **+FTE:** Triage+ based on the ratio of value scores to the total number of FTEs required by the proposal.

- (c) **+Resources**: Triage+ based on the ratio of value scores to a combined measure of the percentage of originally available funding plus the percentage of originally available FTEs.
6. **Triage++ Methods**: Triage++ methods as described in Section 5.2.2 with different CDFs

(a) **Empirically-derived CDF**

- i. **++Cost**: Triage++ based on the ratio of value score to direct dollar cost of the proposal.
- ii. **++FTE**: Triage++ based on the ratio of value scores to the total number of FTEs required by the proposal.
- iii. **++Resources**: Triage++ based on the ratio of value scores to a combined measure of the percentage of originally available funding plus the percentage of originally available FTEs.

(b) **CDF approximated with Triangular distribution**

- i. **++Cost_{tri}**: Triage++ based on the ratio of value score to direct dollar cost of the proposal.
- ii. **++FTE_{tri}**: Triage++ based on the ratio of value scores to the total number of FTEs required by the proposal.
- iii. **++Resources_{tri}**: Triage++ based on the ratio of value scores to a combined measure of the percentage of originally available funding plus the percentage of originally available FTEs.

(c) **CDF approximated via other distributions**

- i. **++Cost_{LN}**: Triage++ based on the ratio of value score to direct dollar cost of the proposal. The required CDF was modeled as a log-normal distribution as described in Section 5.2.4.

- ii. **++FTE_{expo}**: Triage++ based on the ratio of value scores to the total number of FTEs required by the proposal. The required CDF was modeled as an exponential distribution as described in Section 5.2.4.
- iii. **++Resources_{expo}**: Triage++ based on the ratio of value scores to a combined measure of the percentage of originally available funding plus the percentage of originally available FTEs. The required CDF was modeled as an exponential distribution as described in Section 5.2.4.

5.2.4 *Experimental Approach.*

To further explore these methods we constructed a MatLab simulation based on the DP project data. All dollar costs were in \$K, and an initial supply of \$10,000K and 5 FTEs in each of the 18 specialties was available. We simulated project arrivals over 260 days, which corresponds to one year of business days. On the first day of each iteration, seven randomly selected project proposals arrived. On each successive day a number of projects N_{arr} was drawn from a Poisson distribution with λ_{arr} and N_{arr} projects were randomly drawn with replacement from the project pool. On each successive day, a number of departure events N_{dep} was also drawn from a Poisson distribution with λ_{dep} . For the methods which allowed departures, N_{dep} projects were randomly selected for removal from the pool of *pending* projects. Projects that had been *selected* by either the random method or any of the Triage methods were not subject to departure. At the conclusion of the 260 days, any unallocated resources were applied to the pool of pending projects using a binary integer approach to determine the final portfolio.

Additionally, we sought to identify the impact of various aspects of the decision situation on the effectiveness of these methods, as well as the parameters of the methods themselves. The following variables were identified:

1. λ_{arr} The arrival rate of project proposals

2. λ_{dep} The rate at which proposals not selected outright depart
3. Δ_w The sensitivity interval by which the weights in the value hierarchy are allowed to vary in the Triage methods
4. S_{model} The size of the value model, as measured by the number of elements in the value hierarchy
5. α_{model} The α parameter for each of the Triage methods

Finally, as alluded to above, for the methods that required a CDF we used three possible approaches. The first was to use an empirically derived CDF that was constructed in MatLab to precisely match the actual set of 92 project proposals. The second was to use MatLab's Distribution Fitting Tool to fit a triangular distribution to the project proposals. In practice, triangular distributions are often easier to implement when there is little data available about the actual distribution, and we were interested in the loss of value associated with using such a coarse estimate. The third was to normalize the data and use MatLab to provide the best fit among common probability distributions. For the ++Cost method this turned out to be a log-normal, while the ++FTE and ++Resources data were best modeled by exponential distributions. The identification of a log-normal distribution was expected, as it is often a good fit for distributions that are characterized by low, non-negative values with large variance. [65] The identification of an exponential distribution was less expected. The exact fits identified are shown in Table 5.1.

Table 5.1: Distribution Approximations

Method	Shift	Distribution
++Cost	-0.05	log normal(-0.0732, 1.2302)
++FTE	-25	exponential(260.818)
++Resources	-127	exponential(1034)

We employed a full-factorial design with center point for a total of 65 possible variable configurations. Each configuration was run for 100 iterations and the mean portfolio value was captured. We then used JMP software to identify the variables with significant impact on performance. Table 5.2 shows the levels tested.

Table 5.2: Experimental Design

Factor	-1	0	+1
λ_{arr}	0.1	0.2	0.3
λ_{dep}	0.05	0.1	0.15
S_{model}	4	6	8
Δ_W	10%	20%	30%
α_{Triage}	0.25	0.5	0.75
α_{+Cost}	0.4	1.2	2
α_{+FTE}	100	225	350
$\alpha_{+Resources}$	450	975	1500
α_{++} (all methods)	5	10	15

5.3 Factor Effects

While we initially screened all main effects and two-factor interactions it quickly became apparent that main effects alone could explain the preponderance of variability for each of the methods. In no case were more than three main effects identified as significant. In the coded space, JMP identified the factors and factor coefficients shown in Table 5.3.

Several aspects of these results stand out right away. To begin with, we note that in the basic Triage method, the α parameter is not significant. This is a result of the Triage method's failure to consider resource constraints. Selection of high-scoring but resource-intensive projects precluded the later selection of other projects. Thus the overall portfolio

Table 5.3: Factor Effects

Model	Intercept	Variable#1	Variable#2	Variable#3	adj R^2
Hold All	11.73	$-0.76S_{model}$	$+2.42\lambda_{arr}$	—	0.996
Come and Go	8.12	$-0.5S_{model}$	$+3.98\lambda_{arr}$	$-2.05\lambda_{dep}$	0.958
Random	5.91	$-0.4S_{model}$	$+0.72\lambda_{arr}$	—	0.983
Triage	6.86	$-0.66S_{model}$	$+0.97\lambda_{arr}$	$+0.3\Delta_W$	0.719
+Cost	6.19	$-0.47S_{model}$	$+1.11\lambda_{arr}$	$-1.07\alpha_{+Cost}$	0.955
+FTE	7.29	$-0.52S_{model}$	$+1.15\lambda_{arr}$	$-0.2\alpha_{+FTE}$	0.806
+Resources	7.55	$-0.61S_{model}$	$+1.49\lambda_{arr}$	—	0.719
++Cost	6.82	$-0.42S_{model}$	$+0.92\lambda_{arr}$	$+0.57\alpha_{++Cost}$	0.973
++FTE	7.83	$-0.49S_{model}$	$+\lambda_{arr}$	$+0.46\alpha_{++FTE}$	0.955
++Resources	8.41	$-0.51S_{model}$	$+1.36\lambda_{arr}$	$+0.5\alpha_{++Resources}$	0.928
++Cost _{tri}	5.32	$-0.52S_{model}$	$+0.54\lambda_{arr}$	$+1.31\alpha_{++Cost}$	0.97
++FTE _{tri}	5.69	$-0.71S_{model}$	$+0.23\lambda_{arr}$	$+1.5\alpha_{++FTE}$	0.981
++Resources _{tri}	5.55	$-0.68S_{model}$	$+0.27\lambda_{arr}$	$+1.63\alpha_{++Resources}$	0.978
++Cost _{LN}	4.82	$-0.38S_{model}$	—	$+1.44\alpha_{++Cost}$	0.987
++FTE _{expo}	7.86	$-0.63S_{model}$	$+0.98\lambda_{arr}$	$+0.42\alpha_{++FTE}$	0.919
++Resources _{expo}	8.26	$-0.67S_{model}$	$+1.3\lambda_{arr}$	$+0.58\alpha_{++Resources}$	0.906

value was dictated as much by the random arrival order of the projects as by the α parameter employed to screen them. The Triage method was excluded from subsequent analysis.

Next we note the consistent preference for smaller models. Without a true, objective measure of project value external to the model score, it is difficult to draw conclusions from this aspect of the result. A reduced attribute set can make choice problems easier for decision makers by reducing the likelihood of information overload. [32] Smaller models are also preferable in many cases as the cost for obtaining the data necessary to score an alternative is reduced.

At the same time, however, we must consider the possibility that in this case the preference for smaller models is an artifact of the method used to construct the models. Beginning with the full nine-attribute model, the eight-attribute model was developed by removing the lowest weighted attribute and proportionally reallocating its weight to the remaining attributes. The seven-attribute model was then built by removing the lowest

weighted remaining attribute, and so on. The net effect was that the smallest models contained only those attributes that had been weighted the highest to begin with. Since the set of attributes was known to the project submitters when they crafted their proposals, we can assume that they would have focused on these attributes from the outset. By re-allocating weight from the attributes that were likely less pursued to begin with to the attributes that were focused on, it is not surprising that scores should increase. Figure 5.1 shows this behavior.




Figure 5.1: Average Score (95% Confidence Interval) versus Model Size

To further explore this aspect of the problem, we generated a stream of project arrivals and a second stream of departure events. We then provided these streams to models that utilized the same α and Δ_w parameters but different model sizes, in this case four and nine,

and captured their outright selection decisions. The degree of commonality between the decisions made with the two different model sizes is shown in Table 5.4.

Table 5.4: Selection Commonality

Method	Commonality
+Cost	79.38%
+FTE	84.74%
++Cost	79.67%
++Cost _{tri}	93.76%
++Cost _{LN}	87.54%
++FTE	71.88%
++FTE _{tri}	80.47%
++FTE _{expo}	72.34%
++Resources	72.06%
++Resources _{tri}	82.88%
++Resources _{expo}	74.9%

After the model size S_{model} , the next most consistently significant parameter is the arrival rate of the projects λ_{arr} . This is not surprising, as the more projects that arrive within the simulated time frame the greater the chance that low-cost/high-value projects will arrive and positively impact the overall portfolio value. Indeed if there is a surprise with regard to this parameter it is that it was not significant to the ++Cost_{LN} method. Also noted was the relatively small effect this parameter had in the methods that utilized triangular distributions to model the CDFs. This appears to be an artifact of the shape of the triangular distribution and will be discussed later in Section 5.4.3.

Regarding α we note that it is not significant in the +Resources method, which was surprising. We surmise that this was due to the high degree of variability in the +Resources

ratios. While the variance of the +Cost ratios was only 9.34, the variance of the +FTE ratios was a significantly larger 52,138. When the two are then combined to form the +Resources ratio, the resulting variance is an extremely large $8.93 \cdot 10^5$. This calls into question the feasibility of this particular method, at least as applied to this data set. The +Resources method was also excluded from subsequent analysis.

We also note the consistent preference for a less restrictive value of α . In the case of the +Cost and +FTE methods this translates to a lower value, while in the case of the Triage++ methods this means a higher α value. Methods for selecting α values are intended as a future research topic.

5.4 Effectiveness of Extensions

5.4.1 *Impact of Departure Rate.*

We continue our analysis by attempting to understand the relative performance of the methodologies in a “best-case” scenario. To accomplish this, we fixed Δ_W , which was not significant in any of the remaining methods, at its 0 level, 20%, from the full factorial design; we fixed λ_{arr} at its +1 level, 0.3, as all methods had positive coefficients for this factor; we fixed S_{model} at its -1 level, a four attribute model, as all methods had negative coefficients for this factor; we fixed α_{+Cost} and α_{+FTE} at their -1 levels, and all of the α_{++} parameters at their +1 to again coincide with the sign of their coefficients. Finally, we allowed λ_{dep} to vary from a low value of 0.02 up to a high value of 0.3, at which point λ_{dep} and λ_{arr} were equal. We would expect this to only impact the Come and Go method.

Figure 5.2 shows the performance of the two Triage+ methods we pursued, +Cost and +FTE. Further investigation of the +Resources method was abandoned after it was determined the α parameter was not significant in predicting performance, as described in Section 5.3. We first note that, as expected, only the Come and Go method appears sensitive to the departure rate. In the Triage+ methods, project departures only affect those projects that were held pending. These projects are used at the end of the simulation period

to “flesh-out” the portfolio by using any resources not committed to projects that were selected outright.

Next we note that the Come and Go method outperforms the Triage+ method until the departure rate λ_{dep} reaches a value above ≈ 0.20 . The arrival rate λ_{arr} was set at 0.3 for these runs. In the absence of other considerations, e.g. idle resources, this suggests that the Triage+ extensions may only be worth pursuing in an environment where there is a reasonable expectation of a relatively high rate of project departures if they are not selected outright.




Figure 5.2: Performance of Triage+ Methods

We turn now to the Triage++ methods as seen in Figure 5.3. Here we show the performance of the the three Triage++ methods with empirically derived CDFs

(left), triangular approximations to the CDF (center), and log-normal/exponential CDF approximations (right). As expected, varying λ_{dep} only impacted the Come and Go method. We immediately note the relatively poor performance of the triangular approximations to the CDF. This is disappointing, as the triangular distribution is often employed when there is limited data available, as we might expect to be the case in a new continuous decision problem. On the other hand, we note that the log-normal and exponential approximations to the CDFs perform comparably to the empirically fitted ones. This still leaves the not insignificant hurdle of selecting parameters for these distributions with little to no advance knowledge of the actual distributions.




Figure 5.3: Triage++ Performance

Of note in comparing these figures to Figure 5.2 is that the Triage++ methods consistently outperform the Triage+ methods when triangular CDF approximations are not being used. Again we also see that the Come and Go method outperforms the Triage++ methods when $\lambda_{dep} \lesssim \frac{1}{2}\lambda_{arr}$. This is also an improvement over the Triage+ methods where the Come and Go method remained superior while $\lambda_{dep} \lesssim \frac{2}{3}\lambda_{arr}$.

5.4.2 Impact of α .

We now turn our attention to the role played by the value of the α parameter. It is beyond the scope of this paper to investigate methods for effectively selecting a value for

α . For now we limit ourselves to studying the impact of the parameter's value on the overall performance of the Triage methods. In keeping with the method used in Section 5.4.1, we fixed Δ_W , which was not significant in any of the remaining methods, at its 0 level, 20%, from the full factorial design; we fixed λ_{arr} at its +1 level, 0.3, as all methods had positive coefficients for this factor; we fixed S_{model} at its -1 level, a four attribute model, as all methods had negative coefficients for this factor; we fixed λ_{dep} at 0.15, corresponding to $\frac{1}{2}\lambda_{arr}$. We then allowed the value of the α_+ parameter to vary over a range of [0.1, 2.4] for the +Cost method and [25, 475] for the +FTE method. Figure 5.4 shows the resulting performance.

Not surprisingly, neither method achieved portfolio values on par with the Come and Go method, as $\lambda_{dep} \lesssim \frac{2}{3}\lambda_{arr}$. More ominously, we note that the +Cost method performed worse than Random for α_+ values in the latter half of the range. The +FTE method, on the other hand, did not display a huge drop in performance and remained stable, if unspectacular, over a large portion of the α_+ range.




Figure 5.4: Triage+ Performance

Applying a similar approach to the Triage++ methods yielded the results shown in Figure 5.5. The α_{++} values were varied over the range[2,20]. All other parameters were fixed at the levels described above. We again note the significant drop in performance

when utilizing triangular approximations to the CDF. In general, however, performance is much more stable over a significant portion of the α_{++} range. The ++Resources method in particular came close to or exceeded the Come and Go method throughout much of the middle of the range. As noted earlier, this is with λ_{dep} at 0.15, corresponding to $\frac{1}{2}\lambda_{arr}$. Were λ_{dep} to increase, we would expect the ++Resources method more significantly outperform Come and Go.




Figure 5.5: Triage++ Performance

Finally, in Figure 5.6 we see the overhead projection of the surfaces formed by the Come and Go and ++Resources methods as we vary the λ_{dep} and α_{++} parameters. As expected, we see that the two methods behavior is roughly orthogonal: the Come and Go method responds to changes in λ_{dep} while ++Resources responds to changes in α_{++} . The random nature of the project arrival streams leads to the slight texture in the surfaces, though the overall patterns are clear. Come and Go provides decreasing performance as λ_{dep} increases; ++Resources provides improving performance as α_{++} increases, to a point, and then begins to decrease slightly. This is not surprising given the shape of the curves seen in Figure 5.5.




Figure 5.6: ++Resources and Come and Go

5.4.3 Risk.

As we acknowledge from the outset, building a portfolio with any of the Triage extensions is likely to result in an overall loss of total portfolio value as compared to a traditional approach where the selection is made with the complete set of alternatives known. The desire is to gain a sufficient degree of temporal flexibility to offset this loss of value. The exact nature of this tradeoff is unique to every decision maker and situation. We can look at the current decision context however to gain some insight into just how much temporal flexibility we are gaining.

Figure 5.7 shows the commitment of dollars to selected projects as the 260 days of the simulation pass for each of the considered methods, while Figure 5.8 shows the commitment of FTEs. These results were for the “best-case” configurations described in Section 5.4 with λ_{dep} fixed at 0.15, which corresponded to $\frac{1}{2}\lambda_{arr}$.




Figure 5.7: Dollar Commitment versus Time




Figure 5.8: FTE Commitment versus Time

Earlier figures provided a clear picture of the potential loss of portfolio value associated with employing any of the Triage methods in a continuous fashion. Figures 5.7 and 5.8 demonstrate the upside to this risk in terms of temporal flexibility. In a more traditional approach similar to Hold All or Come and Go, no resources would be committed to projects until the end of the decision epoch. The continuous approach allows commitment decisions to take place as projects arrive, and assuming those resources are available, the projects can be worked immediately. Smaller projects may indeed be completed before a commitment decision would even have been made in a traditional approach.

These charts also start to explain the relatively poor performance of the Triage++ methods when using triangular distributions to approximate the required CDFs. All three methods using triangular distributions show a marked tendency to defer commitments until very late in the simulation. Figure 5.9 begins to illuminate why this happens. It shows the empirical CDF for the ++Cost method and the corresponding triangular approximation. The triangular CDF significantly underestimates the probability of a given observation over almost the entire domain. Recall that Triage++ methods calculate $N_{max} = (1 - P(R_{min}(A))) \cdot \lambda \cdot T$ and then select alternatives where $N_{max} \leq \alpha_{++}$. Therefore when $P(R_{min}(A))$ is underestimated, $(1 - P(R_{min}(A)))$ is *overestimated*, and the alternative is less likely to be selected. Only when the remaining time shrinks to a level that offsets this overestimation does the method begin to make significant selections.

5.5 Configuration Guidelines

We set out to begin identifying guidelines to assist decision makers in applying the Triage extensions to continuous decision problems. Understanding that all observations may potentially be particular to this data set, we feel that it was complex enough to effectively represent many real world decision situations. Our experience leads us to the following observations:




Figure 5.9: Comparison of CDF Approximations

5.5.1 *Departure Rate.*

The clearest observation is that in situations where there is a reasonable expectation that alternatives will remain available, deferring the decision as long as possible is preferable to engaging in a continuous decision. This makes intuitive sense, as the more aware the decision maker is of the complete set of alternatives the better. Our experience shows that the Triage+ methods are preferable only when the departure rate for alternatives is roughly $\frac{2}{3}$ or more of the arrival rate. For the Triage++ methods the fraction is closer to $\frac{1}{2}$. In practice of course the decision maker is unlikely to be able to precisely quantify arrival and departure rates. In the case where there is some experience with the decision situation, such as the current iteration of an annual process for example, it is not inconceivable that a reasonable estimate may be formed.

5.5.2 *Model Size.*

All methods investigated performed better with smaller models, i.e., those with fewer hierarchy elements. As described in Section 5.3 the difference in actual score achieved is likely due as much to the process used for developing the different sized models as to any intrinsic superiority of the smaller model. However, as Table 5.4 shows, the actual

project selections are largely the same regardless of model size. Regardless of the decision methodology employed this supports the view that the decision analyst should focus the model on the most relevant set of attributes. Limiting the model to a smaller set of higher valued attributes carries with it the potential benefits of lower cost to acquire scoring data, reduced likelihood of perceived information overload, and an overall process that is faster and easier to implement.

5.5.3 Sensitivity Intervals.

The size of the sensitivity interval used in the Triage methods does not appear to significantly effect the overall performance of the methods. We view this as a positive development as it allows the decision analyst to consider only the decision maker's level of uncertainty in specifying a sensitivity interval without requiring the additional consideration of its impact on the performance of the decision method.

5.5.4 CDF Approximations.

The effectiveness of the Triage++ extensions is largely conditioned on the quality of the CDF approximation used. As we saw, the use of a CDF derived from a triangular distribution performed quite poorly. On the other hand, a CDF approximation using a standard probability distribution fitted to the data performed almost as well as an empirically derived CDF that fit the data exactly. Of course, in a real world decision problem, there will be no data set to fit to, and so the choice of a CDF approximation will be both difficult and fraught.

As we saw in Figure 5.9 the CDF of the +Cost ratios required for the ++Cost method rises very steeply. This was also the case for the other ratios not pictured. Additionally, our previous experience in [41] showed this was the case for value/cost ratios for a purely fictional set of alternatives generated randomly. This suggests distributions that display this general shape may be more useful in modeling CDFs as we move forward.

5.6 Future Work

5.6.1 *Defining α .*

The most obvious avenue for future research on the Triage method lies in developing methods to effectively specify a cutoff value α . This parameter is critical to any of the Triage methods, and as we have seen largely determines the overall effectiveness of the method. To date we have developed no true guidelines to specifying its value however. In both this work and in [41] we have seen that for some α values the extended Triage methods can provide quite effective portfolio generation. We have also seen that other α values can provide truly abysmal performance. To be of real use to decision analysts, we must develop robust methods for specifying α values rather than leaving such a critical piece to chance.

In addition, to date we have tended to utilize fixed values of α . It is not unreasonable to expect that a decision maker might adjust their cutoff value during the course of a continuous decision situation. For example, if relatively few selections are being made throughout the early decision epochs of a continuous decision situation, it might be prudent to lower the value of α . Conversely, if a large number of selections are being made early, it might be prudent to raise the value of α . How best to go about this adjustment is an open question.

5.6.2 *Robust CDF Estimates.*

The Triage++ extensions require an estimate of the CDF of the various ratios in order to function. As we have seen, the quality of this estimate has a significant effect on the quality of the overall decision. In our experiments we utilized a historic dataset that we were able to fit to. A decision analyst facing a new situation will not have this luxury. If possible, the development of guidelines that offer a greater chance of selecting a high quality estimate of the CDF would be of great use in employing these methods.

5.6.3 Value of Information.

An interesting avenue to pursue would be the adaptation of the methods described in [57] to measure the value added by the employment of the extended Triage methods. This would require significant adaptation of the methodology to account for varying project costs and resource constraints, but may provide useful insights into the effectiveness of employing these methods.

VI. The Elicitation and Application of Decision Problem Parameters in the Extended Triage Method

*The contents of this chapter are being prepared for submission to the journal **Decision Analysis**. They have been reformatted to comply with the AFIT style guide.*

The extended Triage method provides an analytically based methodology to evaluate individual alternatives as they arrive over time and engage in a series of accept/reject decisions en route to a final selection of one or more alternatives. The technique uses multi-criteria decision analysis and global sensitivity analysis to compare an alternative's performance to a cutoff value α . We explore two aspects of α in this paper: techniques for eliciting its value from the decision maker, and techniques for allowing this value to vary over time. We also attempt to elicit from the decision maker a predicted cumulative density function that is required for some techniques.

6.1 Introduction

Decision makers in large organizations face a wide variety of decision situations, and are similarly afforded a wide variety of tools and analytical approaches to aid them. This paper is the fourth in a series we have written to address a particular decision environment that is under served in the literature, what we term a *continuous decision problem*. The formal definition of a continuous decision problem is offered in [41]. For our purposes, the key distinction between a continuous and a traditional decision problem lies in the availability of the alternatives. In most, if not all, traditional approaches it is assumed that the entire set of alternatives is known and the methodology seeks to provide the decision maker an objective means to rank order the alternatives, screen out undesirable alternatives, select one or more 'best' alternatives, etc. In a continuous decision problem, this availability of data is not present. Alternatives arrive, asynchronously, over a time

period and the decision maker seeks to assemble the best subset possible within whatever resource constraint set is present. As each alternative arrives, the decision maker has three options:

1. Select the alternative to be part of the final set
2. Reject the alternative
3. Defer a decision on the alternative until a later time period

But how can the decision maker decide which course to take? What analytically sound tools can we offer to assist in this decision? The literature does not offer much on the way of directly applicable approaches. [64] provides a dynamic programming approach to evaluate a stream of alternatives with the goal of selecting a single ‘best’ alternative, while [84] investigates sequential decisions on whether or not to continue a given research and development project. Techniques that are directly applicable to the sequential assembly of a resource constrained portfolio of alternatives arriving over time are lacking.

6.1.1 The Triage Method.

As originally described in [86] the Triage method is not a selection method, but a screening one. It is designed as the first stage in a two-step project evaluation process. The goal of the Triage method is to evaluate individual alternatives as they present themselves and determine whether or not the alternative is worth further, more detailed evaluation. In this method, each alternative is first scored using a linear additive value model. A linear program (LP) is then used to determine the set of weights within a given sensitivity interval defined by the decision maker that will maximize the alternative’s score. As shown in [86] the symmetrical nature of the linear additive value model ensures that if an alternative A achieves a value score $V(A)$ in the original model and a maximum score $V_{max}(A) = V(A) + \Delta(A)$ under the LP-derived weights, then there is no need to re-run

the LP to minimize the alternative's score. The worst the alternative can score within the sensitivity interval is $V_{min}(A) = V(A) - \Delta(A)$.

These values, $V_{min}(A)$ and $V_{max}(A)$, are now compared to the cutoff parameter $\alpha \in [0, 1]$ and the alternative is triaged into one of three categories:

1. The alternative is selected to continue in the evaluation process if $V_{min}(A) > \alpha$
2. The alternative is rejected if $V_{max}(A) \leq \alpha$
3. The alternative is help pending otherwise, and may be cycled back through the process at a later date if the decision maker modifies α

This evaluation takes place without regard to the resource requirements of the individual alternatives. While appropriate for a screening model, this omission makes the approach infeasible for selecting alternatives. If the first alternative with $V_{min} > \alpha$ has resource requirements that consume the entire available budget, it would still be selected despite the fact that this is likely not an efficient allocation of resources. What is required is a way to extend the Triage method to consider resource requirements.

6.1.2 The Extended Triage Method.

In [42] we propose two fundamental approaches to extend the Triage method and add this resource consideration. The first, which we term Triage+, involves taking the ratio of $V_{max}(A)$ and $V_{min}(A)$ to a critical resource cost, C_{crit} , of the alternative. If there is a single resource constraint then the construction of this ratio is straightforward. If there are multiple resources involved, then C_{crit} can be specified as the one considered most likely to be binding, or constructed as a combination of the various resource requirements. These ratios, $R_{max} = \frac{V_{max}(A)}{C_{crit}}$ and $R_{min}(A) = \frac{V_{min}(A)}{C_{crit}}$ are then compared to a new cutoff value α_+ and triaged as in the original method. This comparison discourages the selection of alternatives that are resource intensive. The selection of a value for α_+ is not necessarily straightforward however as it is no longer bound to the range $[0,1]$.

The second extension, termed Triage++, requires more familiarity with the decision environment, and is best suited for new iterations of a previously engaged cyclical process. It requires two pieces of information beyond the Triage+ method, an estimate of the cumulative distribution function (CDF) of the ratio $\frac{V(A)}{C_{crit}}$ and an estimate of the arrival rate λ_{arr} of new alternatives. As each alternative arrives then, it is scored and the ratios R_{max} and R_{min} are formed as in the Triage+ method. These are then combined with the estimated CDF and the arrival rate λ_{arr} to arrive at an estimate for the number of alternatives the decision maker can expect to see in the remaining time whose benefit/constraint ratio exceeds R_{max} and R_{min} and these numbers are compared to a cutoff value α_{++} . The alternative is then triaged based on the outcome of this comparison.

6.1.3 Key Issues.

Our earlier work in [41] first introduced the Triage extensions and applied them to a sample problem, and then investigated the effects of decision problem parameters on their effectiveness when applied to a more realistic data set in [42]. We turn our attention now to two aspects of application that have not yet been investigated: how to elicit critical pieces of information from the decision maker, namely the value of the α parameter and an estimate of the required CDF, and expanded options for how to apply α once its value has been elicited. In our work to date we have not elicited an α parameter from a decision maker, but have instead applied a range of values to gauge effectiveness. Additionally, for each decision modeled we have used a constant value for α throughout the decision timeline. We will now examine methods for eliciting α as well as various methods for allowing α to vary over the decision timeline. With regard to CDF, to date we have only used either an empirically derived CDF, or a distribution (log-normal, triangular, etc.) fitted to the empirical data. This will mark our first effort at estimating this part of the problem.

The remainder of this paper is organized as follows: Section 6.2 provides background on the dataset employed in this effort, the techniques used to elicit critical values from the

decision maker, and the manner in which those values will be applied. Section 6.3 provides the outcome of our elicitation efforts and the results the extended Triage techniques achieve using them. Finally, Section 6.4 summarizes our conclusions.

6.2 Background

6.2.1 *The Development Planning Dataset.*

We continue with the dataset employed in [42]. This data is drawn from the HQ Headquarters (HQ) United States Air Force Materiel Command’s (AFMC) Development Planning (DP) activity. From Fiscal Year (FY) 2011 to FY 2015 HQ AFMC conducted an annual call for DP project proposals. Each proposal included data on the project to be pursued, as well as the level of funding and manpower resources it would require. These were collected, scored against a value model, and an LP was then utilized to determine the subset with the highest total value score that could be accommodated within funding constraints and available manpower, expressed in full time equivalent (FTE) positions in 19 different specialties. We have collected the value model scores and resource requirements for a pool of 92 unique DP project proposals collected during that time period.

For the simulations conducted in this paper we draw from this pool, with replacement, with new arrivals being Poisson distributed with an average arrival rate $\lambda_{arr} = 0.3$. This corresponds to an average of 78 proposals arriving per simulated year. Because the actual DP process employed a single, annual data call we do not have a historical record of temporally distributed arrivals, and thus substitute the Poisson arrival process. In our previous work we allowed the value of λ_{arr} to vary and, as expected, the higher the arrival rate the better the total value of the portfolios achieved, as there are more options to choose from and a higher likelihood of high value projects. The value of 0.3 used in this paper corresponds to the upper bound of the range examined in [42]. We also assume that alternatives not selected outright may cease to be available before the end of the decision timeline. Each day a number of ‘departure events’ is drawn from a Poisson distribution

with $\lambda_{dep}=0.15$. For each departure event a project proposal was randomly selected from the pool of projects being held pending and removed. Again, there is no historical basis for treatment of project departures as the original DP selection process involved a single decision event. The substitution of a Poisson process with $\lambda_{dep} \leq \lambda_{arr}$ is a reasonable approach.

We do not seek to duplicate the historical experience of the DP prioritization process, nor can we compare our selected portfolios directly to the historical record. As mentioned above, there is no historical analog to alternatives arriving over time or departing due to the annual data call. More importantly, in the actual DP prioritization process, there was an additional step after the LP selection of an optimal portfolio, termed “professional military judgment” (PMJ). During the PMJ phase, sponsoring agencies were allowed to review the list of selected projects. If a project was not selected, but was highly desired by the sponsoring agency, they were allowed to “de-scope” one or more of their selected projects and apply the freed resources to non-selected projects, so long as the total level of resources committed to that agency did not increase. Thus alternative X that appeared in the final selection list may or may not correspond directly to the alternative X whose score data was captured. In effect, the process allowed sponsoring agencies to re-define the alternatives and these new alternatives were never scored.

6.2.2 Elicitation Methods.

A key step in developing any decision model is eliciting preference information from the decision maker. Use of the extended Triage method involves eliciting even more information, which can prove difficult. This is an important step not only because the elicited values directly impact the performance of the method, but also because they can impact the level of confidence the decision maker has in the results regardless of their actual performance [3].

6.2.2.1 α Value Elicitation.

We adapt approximation techniques for probability distributions, such as those in [51] and [22], to our effort to elicit α values from the decision maker. These techniques call upon the decision maker to provide numerical inputs to specific questions, and thus avoid issues with verbal expressions of probability [12]. For this effort we consulted a member of the AFMC staff who was familiar with the DP prioritization effort and with the development planning process in general, but who had no direct experience with the process during the time period corresponding to our sample data. This helped ensure that his views were not colored by any past experience which might allow values to be elicited that were more accurate than might otherwise be possible.

We first made sure our surrogate decision maker was familiar with the structure of the model. In practice this step would be unnecessary as the decision maker, having just been involved in the construction of the model, would be intimately familiar with its elements. We then postulated the existence of three fictional alternatives with overall value scores of 0.9, 0.1, and 0.5 in that order. No detail as to how the alternatives scored on their individual attributes was provided, simply an overall score. We asked the decision maker to specify what levels of resources he would be willing to apply to these alternatives.

It would be unrealistic to expect the decision maker to provide a detailed answer given the abstract nature of the fictional alternatives. As mentioned above, each DP proposal was associated with a direct dollar cost and manpower requirements in 19 different technical specialties. Given an abstract alternative with no information on the technical domain it is associated with, it is impossible to specify an exact set of resources by discipline. Instead, we grouped the 19 specialties into nine by disregarding geographic classifications. For example, three of the 19 specialties are Program Managers at each of three different product centers. In the reduced set, we simply spoke to “Program Managers” without regard to

location. We then asked the decision maker to provide a maximum level of resources in each of these nine categories.

This led directly to a further issue. The goal of this exercise was to arrive at ratios of value score to resource requirements that could be used as α values in the simulation. Our decision maker had provided us the maximum level of resources he would commit for each of the nine broad categories. In practice however, DP project submissions typically required manpower from only three or four of the nine manpower specialties. Our surrogate decision maker, having no direct experience with the actual historical data, could not have known this. An actual participant facing a first iteration of a continuous decision problem similarly would be hard pressed to predict the type of resource requirements forthcoming alternatives might require. We therefore formed ratios based on the entire set of resources he had identified, as to do otherwise would have taken advantage of information that would not be present in a realistic application.

We then asked the decision maker to consider the nine attributes that were scored in the value model, and to then specify what he considered a “minimally acceptable alternative.” The guidance was to specify a set of attribute scores for an acceptable alternative such that, should any of them decrease, the alternative would no longer be considered for acceptance at all regardless of its resource requirements. While we offered the decision maker the opportunity to specify multiple combinations, he thought it best to only specify one. At this point, we did not appraise the decision maker of what overall score such an alternative would achieve. Once specified, we asked the decision maker to again specify resources he would be willing to commit to such a project.

6.2.2.2 CDF Elicitation.

Making predictions about probability distributions is a difficult task, particularly when you are reliant upon a single decision maker with no historical baseline to refer to [8]. We chose to directly apply techniques such as those cited in Section 6.2.2.1 to obtain

the decision maker's estimate of the CDF for the value to resource cost ratio. We again described to the decision maker that we were considering the ratio of an alternative's value score to its resource cost, as described in Equation (6.1). We then asked him to specify the value on the X-axis that he thought would correspond to a number of specified cumulative probabilities on the Y-axis. The choice of Y-values to ask was governed by the methods employed, which were the extended Pearson-Tukey (PT) and extended Swanson-Megill (SM) approximations described in [51] and the bracket median (BM) approach described in [22]. In total this meant asking for X-axis values corresponding to cumulative probabilities of 0.05, 0.10, 0.3, 0.5, 0.7, 0.9, and 0.95. In addition to the points specified by those methodologies, we asked the decision maker for X-axis values corresponding to predicted probabilities of 0 and 1. Given those responses, we then fit smooth curves through the resulting points to arrive at our estimates for the CDF.

$$R(A) = \frac{V(A)}{\frac{\sum_{i=1}^{19} FTE_i(A)}{\sum_{i=1}^{19} Initial(FTE_i)} + \frac{Cost(A)}{Initial(Budget)}} \quad (6.1)$$

6.2.3 *Application Methods.*

There are a variety of ways the α parameter can be employed in any of the Triage methods. The most direct is to simply fix the value of α for the duration of the decision process. This was the method employed in [41] and [42]. In [86] α was varied on a monthly basis, linearly decreasing from an upper bound to a lower one. In this work we will further explore methods for allowing α to vary. Our aim is not to exhaustively identify methods, for they can be uniquely tailored to the decision situation, nor to identify a single 'best' approach, but rather to gain some initial insight into the relative effectiveness of general classes of application methodology.

6.2.3.1 Varying α on a Schedule.

We should note that varying the value of α has different implications for the Triage+ and Triage++ methods. In the case of Triage+ as the ratio of value to resource cost required for selection increases, this implies that alternatives must either score higher or achieve their score at a lower resource cost. In short, the selection of alternatives becomes more restrictive. In the case of the Triage++ approach, α serves as a cutoff expressed in the number of “better” alternatives the decision maker should expect to see, and they select the current alternative if the expected number is less than α . Thus as α increases the process is effectively becoming less restrictive about the selected alternatives.

Functions of an exponential form are widely used in forecasting, as in [71], and decision making, as in [68]. [59] utilizes an exponential form for the specification of single dimension value functions, as shown in Equations (6.2) and (6.3) for increasing and decreasing functions in the range [Low, High] respectively. The parameter ρ controls the shape of the curve between the Low and High bounds.

$$v(x) = \begin{cases} \frac{1 - e^{\frac{-(x-\text{Low})}{\rho}}}{1 - e^{\frac{-(\text{High}-\text{Low})}{\rho}}}, & \rho \neq \infty \\ \frac{x}{\text{High} - \text{Low}}, & \rho = \infty \end{cases} \quad (6.2)$$

$$v(x) = \begin{cases} \frac{1 - e^{\frac{-(\text{High}-x)}{\rho}}}{1 - e^{\frac{-(\text{High}-\text{Low})}{\rho}}}, & \rho \neq \infty \\ \frac{\text{High} - x}{\text{High} - \text{Low}}, & \rho = \infty \end{cases} \quad (6.3)$$

We will utilize this form to smoothly vary our α values over a pre-defined range, utilizing a variety of values for the ρ parameter to control how aggressively the value of α is modified. Figure 6.1 shows an example of a monotonically increasing α parameter varied over 260 days in the range [450,1500] for a variety of values of ρ . We use a ρ value of 0 to represent the case where $\rho = \infty$, resulting in a linear progression between the lower and upper bounds.




Figure 6.1: α Values for Selected Values of ρ

Utilizing Equations (6.2) and (6.3) we can investigate the performance of the extended Triage methods using α values that either increase or decrease over time at various rates governed by ρ .

6.2.3.2 Varying α Reactively.

In this method we once again employ an exponential form. In this case however, we use it to provide a target level of resource allocation as a function of time. The ρ parameter again controls the shape of the curve, and thus how aggressively the decision maker desires to commit their resources. At the beginning of each simulation day, the curve is consulted to determine whether the target level of resource commitment has been achieved by the selections made to date. If the level achieved is below the target, α is modified to be less restrictive by an increment of $\gamma\%$. Similarly if the level achieved is running ahead of the target then α is made more restrictive by an increment of $\gamma\%$. This removes the restriction that α vary monotonically. In our investigation we adjust α daily, but there is nothing to mandate this particular period. We also explore a variety of values for the size of the adjustment increment γ .

6.2.3.3 Varying the Calculation of R.

As alluded to in Section 6.1.2 when forming the ratios R_{min} and R_{max} the denominator is a measure of the level of resources required by the alternative. In cases where there is more than a single resource constraint some method of combining resource requirements is required. In our work to date, resource requirements have been of two types, a direct dollar cost and a number of FTE's required in one or more of 19 different specialties. We have approached consideration of these requirements via the Triage+ and Triage++ methods by forming the required ratios in three ways:

1. Consideration of FTEs only (+FTE/++FTE) — $R(A) = \frac{V(A)}{\sum_{i=1}^{19} FTE_i(A)}$
2. Consideration of direct cost only (+Cost/++Cost)— $R(A) = \frac{V(A)}{Cost(A)}$
3. Combined FTE and direct cost (+Resources/++Resources) — $R(A) = \frac{V(A)}{\frac{\sum_{i=1}^{19} FTE_i(A)}{\sum_{i=1}^{19} FTE_i} + \frac{Cost(A)}{\text{Initial Budget}}}$

Note that in the third method, we calculate the percentage of the *original* level of funding and FTEs that are required by the alternative under consideration. In this work, we investigate modifying this calculation to consider the percentage of *remaining* funding and FTEs the alternative requires, as shown in Equation (6.4). As dollars and manpower are committed to selected projects, the pool of remaining resources grows smaller, and thus the percentage of available resources required will increase. By fixing the value of α but modifying the ratio calculation in this manner, we essentially impose a monotonically increasing α but at a variable rate.

$$R(A) = \frac{V(A)}{\frac{\sum_{i=1}^{19} FTE_i(A)}{\sum_{i=1}^{19} \text{Remaining}(FTE_i)} + \frac{Cost(A)}{\text{Remaining}(\text{Budget})}} \quad (6.4)$$

6.2.3.4 *Traditional Methods.*

As a basis for comparison, we model three more traditional approaches to the decision situation.

1. Hold All — all arriving projects are held with no possibility of project departure. At the end of the decision period the best scoring portfolio is assembled. This provides an upper bound on the value achievable for the given stream of project arrivals.
2. Come and Go — similar to the Hold All approach, except that projects can depart before the portfolio is constructed at the end of the decision period.
3. Random — each arriving project that is feasible in terms of the current level of resource constraints has a 50% chance of being selected for the final portfolio and a 50% chance of being rejected and no longer considered.

6.3 Results

6.3.1 *Elicitation Results.*

6.3.1.1 α Elicitation.

As described in Section 6.2.2.1 we began by attempting to determine values to use for the cutoff parameter α . Our first exercise was to present the decision maker with fictional alternatives that scored 0.9, 0.1, and 0.5 and then ask what resource levels he would commit to those projects, in terms of FTEs and dollars. The results of this exercise and the resulting ratios are shown in Table 6.1.

Table 6.1: Resources Applied to Fictional Alternatives

Alternative Score	Program Mgr	Engineer	Financial Mgr	Contracting	Logistics	Intelligence	Systems Engineer	Cost Estimator	Modeling & Simulation	Cost (\$K)	Ratio
0.9	1	1	0.5	0.5	0.5	1	1	0.5	0.5	150	4120
0.5	0.625	0.665	0.415	0.375	0.3	1	0.665	0.415	0.375	108	3140
0.1	0.25	0.33	0.33	0.25	0.1	1	0.33	0.33	0.25	67	1000

We then asked the decision maker to specify a “minimally acceptable” alternative and to specify the resources he would be willing to commit to it. The resulting alternative achieved a value score of 0.477, although the decision maker was not aware of this calculation when he chose resource levels. Table 6.2 shows the results.

Table 6.2: Resources Applied to Minimally Acceptable Alternative

Alternative Score	Program Mgr	Engineer	Financial Mgr	Contracting	Logistics	Intelligence	Systems Engineer	Cost Estimator	Modeling & Simulation	Cost (\$K)	Ratio
0.477	0.25	0.33	0.33	0.25	0.1	1	0.33	0.33	0.25	67	4120

The decision maker noted that he greatly preferred the minimally acceptable alternative approach to the fictional alternative approach. This preference could prove important, as [3] showed that the choice of elicitation technique can significantly impact the level of confidence decision makers show in the resulting model products. The differing approaches also produced substantially different results. As mentioned, the decision maker only knew the levels achieved for the various model attributed when specifying his minimally acceptable alternative, not the resulting score such an alternative would achieve. Although the 0.477 score that the minimally acceptable alternative received was very close to the fictional alternative scoring 0.5, the decision maker was willing to commit significantly more resources to the 0.5 alternative. In fact, the decision maker’s initial inclination was to resource the minimally acceptable alternative at the same level as the fictional alternative scoring 0.1. He reconsidered however and specified that values shown in Table 6.2 which still reflected a pessimistic assessment of the value the alternative would ultimately achieve. Afterward, he remarked that he had expected the minimally acceptable alternative to have scored in a range much closer to 0.25.

Following this exercise, the minimum value for α_+ arrived at was 1000, and the maximum was 4120, as can be seen in Tables 6.2 and 6.1. For the purposes of our simulations, we chose to use these to define the range of varying α_+ values. For those

methods where a fixed value of α_+ is employed, the simulation will be run across the range to investigate performance. Discussions with the decision maker led to a similarly applied range of α_{++} values of [5,15]. This was a far more ‘intuitive’ decision on the part of the decision maker, arrived at after a brief discussion on the Triage++ method.

6.3.1.2 CDF Elicitation.

As described in Section 6.2.2.2 we utilized three approaches to arrive at different estimates of the CDF required by the Triage++ methods. Our expectation was that this would be a very difficult task for a decision maker to approach with no basis in historical data to guide the exercise. The results showed that it was even more difficult than we had imagined. Table 6.3 shows the values chosen by the decision maker via the different methods as well as the corresponding empirical data points, while Figure 6.2 shows the resulting curves. It is clear that the elicited CDFs correspond very poorly to the actual distribution of the data. Again, this is not surprising, but confirms the view that the Triage++ techniques are best suited to decision situations in which there is historical data to draw upon.

Table 6.3: Elicited CDF Points

Method	p(X)=0	p(X)=0.05	p(X)=0.1	p(X)=0.3	p(X)=0.5	p(X)=0.7	p(X)=0.9	p(X)=0.95	p(X)=1
Extended Pearson-Tukey	0	1000			3000			5500	10000
Extended Swanson-Megill	0		1750		3000		5250		10000
Bracket Median	0		1750	2500	3000	3750	5250		10000
Empirical Values	165	234	286	538	865	1243	2129	3102	5080

6.3.2 Simulation Results.

6.3.2.1 Fixed α .

As described in Section 6.3.1.1 we used α_+ values in the range [1000, 4120] for the Triage+ method, and α_{++} values in the range [5,15] for the Triage++ method. We began




Figure 6.2: Elicited CDFs

our investigation by simply using fixed values for α . Figure 6.3 shows the performance for various α values utilizing this approach.




Figure 6.3: Fixed α Values

We first note that the +Resources method is able to approach the performance of the traditional Come and Go method for a very small range of α_+ values at the lower end of the investigated range. This is encouraging given the difficulties already described in selecting

α_+ values without the benefit of a historical baseline to refer to. In or previous work, we had typically seen the +Resources method approach Come and Go in cases where $\lambda_{dep} \gtrsim \frac{2}{3}\lambda_{arr}$. Recall that for this simulation $\lambda_{dep} = \frac{1}{2}\lambda_{arr}$.

As the center chart shows, the ++Resources method can match the traditional Come and Go approach while offering the decision maker added temporal flexibility. This is the case for a large portion of the investigated range of α_{++} values. However, this benefits from access to the empirical CDF derived from historical data. As we see in the rightmost chart, performance using the estimated CDF was abysmal, barely approaching that of the Random method at the extreme end of the investigated α_{++} values. This will prove to be a recurring theme.

6.3.2.2 α Varied by Schedule.

As described in Section 6.2.3 we utilize an exponential form shown in Equations (6.5) and (6.6) as the basis for defining our α schedule for increasing and decreasing values of α respectively. First described in [59] for use in the definition of single dimension value functions, we use them to define the value of α as a function of time, $t \in [0, 260]$ days.

$$\alpha = \begin{cases} \frac{1 - e^{\frac{-t}{\rho}}}{1 - e^{\frac{-260}{\rho}}}, & \rho \neq \infty \\ \frac{t}{260}, & \rho = \infty \end{cases} \quad (6.5)$$

$$\alpha = \begin{cases} \frac{1 - e^{\frac{-(260-t)}{\rho}}}{1 - e^{\frac{-260}{\rho}}}, & \rho \neq \infty \\ \frac{260 - t}{260}, & \rho = \infty \end{cases} \quad (6.6)$$

As shown, these equations take on values $\in [0, 1]$. By combining them with a minimum value and a scale factor we can use them to define α values in a desired range. For example, Equation (6.7) shows the form used to define an increasing value of $\alpha_+ \in [450, 1050]$. The resulting values of α_{++} for various levels of the ρ parameter were shown in Figure 6.1.

$$\alpha_+ = \begin{cases} 450 + 1050 \cdot \frac{1 - e^{\frac{-t}{\rho}}}{1 - e^{\frac{-260}{\rho}}}, & \rho \neq \infty \\ 450 + 1050 \cdot \frac{t}{260}, & \rho = \infty \end{cases} \quad (6.7)$$

In Figure 6.4 we see the overall performance of the increasing α method for various values of the ρ parameter. For the +Resources method, we begin α_+ at a value of 1000 and allow it to increase to a value of 4120 at a rate governed by the value of ρ . For the ++Resources method the α_{++} values vary from 5 to 15.




Figure 6.4: Increasing α by Schedule

We first note that the +Resources method is again able to approach the performance of the traditional Come and Go method for a very small range of ρ values that are slightly negative. Performance for ρ values either side of this range was universally poor.

In the center figure we see the performance of the ++Resources method when used with the empirically derived CDF. For a small range of slightly positive values of ρ we see that this method was capable of matching the portfolio value performance of the Come and Go method. In practice, the decision maker would also gain the temporal flexibility of being able to make immediate resource commitments. However, the figure on the right shows the same method when used with the estimated CDFs. They all cluster together, which is

not surprising as Figure 6.2 showed they were not significantly different from one another, with performance that at best matches that of the Random method.

Figure 6.5 shows the corresponding performance for an α value that starts at the upper bound and decreases as time proceeds. We see similar performance, albeit reflected around 0. The peak achieved by the +Resources method is not quite as high as with an increasing α_+ , while the ++Resources method using the empirical CDF achieves values comparable to the Come and Go method over a wider range of ρ values. The ++Resources methods using the elicited CDFs maintain their poor performance.




Figure 6.5: Decreasing α by Schedule

6.3.2.3 α Varied Reactively.

Here we use the exponential form to define a target level of resource commitment rather than the value of α . At the start of each simulation day we compare the current level of commitment to the target level and adjust the value of α accordingly, as described in Section 6.2.3.2. We measure resource commitment in terms of the percentage of the originally available resources committed, so we can use the exponential form shown in Equation (6.2) to smoothly transition the allocation target through values in the range $[0,1]$.

As applied in this work, the current level of resource allocation was strictly compared to the target. Since it is highly unlikely they would match exactly, the value of α was

modified practically every day. In practice, the user would likely specify an interval and only modify the value of α if the discrepancy between the current and target allocation levels exceeded it.

Figure 6.6 shows the results for the +Resources method when applied to a reactively varying α_+ . The figures on the left represent the case where α_+ begins at its lower bound, while those on the right are where α_+ began at its upper bound. The top figures show overall portfolio value, while the contour plots on the bottom represent the difference between the portfolio value achieved by the +Resources method and that achieved by the Come and Go method with the same project stream.




Figure 6.6: +Resources Reactive

In the case where α_+ began at its low bound, portfolio performance came close to that of the Come and Go method, though it did not match it. The differences were in the range

[-1.65, -0.39] with an average delta of -0.97. There is no clear pattern to the performance or deviations, though in general performance was slightly better for negative ρ values.

The case where α_+ began at its high bound, i.e. at its most restrictive, did not perform nearly as well. Though the surface plot in the upper right shows a large jump in portfolio values for a certain range of parameters, it still does not approach the performance of the Come and Go method. The differences for this method were in the range [-9.73, -7.91] with an average value of -9.02. Positive values of ρ and large values of γ were associated with the uptick in performance, but the value achieved was still not comparable to traditional methods.

Figure 6.7 shows the results for the ++Resources method when applied to a reactively varying α_{++} with the Pearson-Tukey CDF approximation. Of the three CDF approximations this was the one that returned the highest average portfolio value, though it still did not perform very well. Again, the figures on the left represent the case where α_{++} begins at its lower bound, while those on the right are where α_{++} began at its upper bound. The top figures show overall portfolio value, while the contour plots on the bottom represent the difference between the portfolio value achieved by the ++Resources method and that achieved by the Come and Go method with the same project stream.

Overall performance in this methodology was disappointing, though this was not surprising given the previous findings on the accuracy of the estimated CDF. For the case where α_{++} began at its low bound, the delta from Come and Go was in the range [-7.7, -6.62] with an average value of -7.07. For the case where α_{++} began at its high bound, the deltas were in the range [-4.7, -2.99] with an average value of -3.88

Figure 6.8 shows similar results when the empirically derived CDF is used rather than the Pearson-Tukey estimate. With α_{++} beginning at its low bound the deltas from Come and Go are in the range [-2.95, -1.18] with an average of -2.05. With α_{++} starting at its high bound, the performance is considerably better with deltas in the range [-0.54, 0.35] with an




Figure 6.7: ++Resources Reactive with Pearson-Tukey CDF Approximation

average value of -0.11. This serves to reinforce our view that the ++Resources method can perform well when historical data is available to approximate the CDF, but is likely not a viable option in the absence of such data.

6.3.2.4 New Calculation of R .

We turn now to our consideration of a modified method for calculating the ratio R used in both of the extended Triage methods. As described in Section 6.2.3.3 we investigate calculating R as a ratio of value score to *remaining* resources rather than to the beginning resource levels. As projects are selected outright and their resource requirements are subtracted from the starting levels, this modification will result in the size of the denominator increasing and thus the overall value of the ratio R decreasing.




Figure 6.8: ++Resources Reactive with Empirical CDF

The impact of this change is difficult to predict. The value of the ratio R for a given alternative now depends not only on the intrinsic properties of the alternative, but also on the order in which it arrived. This had two important implications for our proposed extensions. First, α_+ values chosen for the Triage+ method that are appropriate for the original R calculation are not necessarily appropriate for the new calculation. We did not, however, attempt to elicit new values as it is difficult enough for a decision maker to arrive at α_+ values with the traditional calculation. Second, the CDF that the Triage++ extension relies upon is no longer valid, as again the value of R for a given alternative varies depending on its arrival order. Again, we did not attempt to modify the CDFs used to date, but rather simply observe the difference in performance.

Figure 6.9 shows the results for fixed values of α . The top left graph shows performance for the range of α_+ values investigated to date. We note that the new methodology's highest value is at the lowest α_+ value considered. Given that the modified calculation is likely to drive down the values of R , we decreased the lower bound of α_+ values from 1000 to 100, and the results are shown in the top right graph. We note that the new R calculation method does peak at a lower value of α_+ , but still does not achieve the performance of the Come and Go method.

The lower two charts in Figure 6.9 show the results for the Triage++ method using the empirical and Pearson-Tukey elicited CDFs. Only the Pearson-Tukey elicited results are shown, as it was the best performing of the three elicited CDFs, though the difference was slight. Using the empirical CDF we note little difference in performance. Similarly, performance is virtually unchanged, and again poor, using the elicited CDF.




Figure 6.9: Fixed α Values with New **R** Calculation

Figure 6.10 shows six graphs for α values varied by schedule using the new calculation of R . Across the top row we see results with an increasing value of α first for the Triage+ method, then the Triage++ method with the empirically derived CDF, and finally the Triage++ method with the Swanson-Megill method. Of the three elicited CDFs this was the best performer in this methodology, though again the difference was very slight. The three figures on the second row show the corresponding results for a decreasing value of α .

In general we note that the original calculation of R provides better performance in the Triage+ method regardless of the direction of variation. In the Triage++ method utilizing the empirical CDF the choice of R calculation makes little difference, with the only significant deviation being that it performs slightly worse then the original calculation for positive values of ρ and decreasing values of α_{++} . Performance is indiscernible, and poor, utilizing the elicited CDF.




Figure 6.10: α Values Varied by Schedule with New **R** Calculation

In Figure 6.11 we see several aspects of the results for the use of the modified R calculation with the Triage+ method and the reactive setting of the α_+ value. Along the top row we see results when α_+ began at its low bound, while the bottom row shows results for beginning at the high bound. The first surface plot shows overall portfolio value achieved, while the two contour plots show the difference between the portfolio value achieved by the Triage+ method and that achieved by the Come and Go method, and by the Triage+ method with the original R calculation respectively.

There are very few cases when α_+ began at its low bound where the modified calculation of R provides better performance than the original, and then only slightly. The differences lie in the range $[-1.9, 0.6]$ with an average value of -0.67 . Beyond this, the use of the modified R calculation provides universally lower performance than the original. As with the original calculation of R , beginning α_+ at its high bound provides very poor performance.

Finally, in Figure 6.12 we see a similar display of the performance of the Triage++ method when used with the empirical CDF. As noted in Section 6.3.2.4 this CDF is not necessarily appropriate as the distribution of R values is no longer stationary in this formulation. Surprisingly however, this offered the one occasion when the modified calculation of R seemed to offer some benefit.

On the top row we see results when the initial value of α_{++} was set at the low bound. Nowhere did this method exceed either the performance of the Come and Go method or that of the original R calculation. On the bottom row however we see that beginning α_{++} at the high bound did offer some improvement in performance. When compared to the Come and Go method, the differences were in the range $[-0.5, 0.88]$ with an average value of 0.23 . Compared to the original calculation of R the differences were in the range $[-0.26, 1.03]$ with an average value of 0.32 .




Figure 6.11: +Resources Reactive with New $\mathbf{R}$ Calculation

We omit display of the Triage++ method when used with the elicited CDFs as it continues the pattern of exhibiting very poor performance. Average portfolio values were on the order of 4.4 points below the Come and Go method, a difference of approximately 43%.




Figure 6.12: ++Resources Reactive with New $\mathbf{R}$ Calculation

6.4 Conclusions

While we recognize that there is a limited capacity to draw broad conclusions from the study of a single decision problem such as ours, we offer the following observations and conclusions.

6.4.1 *The Need for Historical Data.*

The Triage++ methodology relies heavily upon the use of the CDF for the value to cost ratio R . We have shown that the method can be highly effective when used with an empirically derived CDF, or with a functional form such as the log-normal fitted to empirical data. What is important is that the utilized CDF provide a quality approximation of the actual CDF. In [42] we showed that the use of a triangular distribution, even when fitted to historical data, provided very poor performance. This works reinforces the idea that a quality CDF approximation is vital. Even when using elicitation techniques that move beyond the triangular distribution, we found universally poor performance with elicited CDF estimates.

Consideration of a ratio of value score, derived from a multi-criteria model, to resource cost, as expressed in a combination of various factors, is not an intuitive exercise, nor is it one that decision makers are likely to have engaged in previously. Combine that with the inherent uncertainty of a decision problem in which the field of alternatives is unknown, and it is not surprising that even an educated estimate of the CDF for these ratios is likely to be highly inaccurate. As we have seen though, this has a highly negative effect on the performance of the methodology, with overall portfolio values dropping 60% or more from the values achieved with the empirical CDF. The difficulty in estimating this critical parameter combined with the consequences of a poor estimate argue strongly that it should only be considered for newer iterations of a process where there is a historical record that can confidently be used to characterize the expected alternative stream.

6.4.2 *The Modified Calculation of R .*

While there was an intuitive appeal to calculating R as a ratio of value to *remaining* resources as opposed to the originally available resource levels, it was not borne out by the simulation results. In only one small range of parameter settings within one application method did we see any improvement in the performance achieved by the modified R

calculation, as described in Section 6.3.2.4. Given that this improvement occurred in a Triage++ application, where the CDF approximation being used was based on the original R calculation, we cannot be confident at all that this is reflective of an actual methodological improvement and not just a fortunate stream of alternatives. Further investigation is warranted as it would seem that basing decisions on current resource levels, as opposed to starting ones, would be a more effective use of all available information.

6.4.3 Varying α .

Table 6.4 provides an overview of the various methods used to apply the α parameter ordered both by average and maximum achieved portfolio value. We first note that in general the Triage++ methodologies were the highest scoring, but with the caveat that the values displayed are for the method using the empirical CDF. Those achieved with the elicited CDFs would appear at the bottom of the list.

Table 6.4: Performance Summary

Method			Average
Triage++	Reactive	α_{++} Begins High	10.21
Triage++	Schedule	α_{++} Begins High	10.05
Triage++	Fixed α_{++}		9.81
Triage++	Schedule	α_{++} Begins Low	9.46
Triage+	Reactive	α_+ Begins Low	9.35
Triage++	Reactive	α_{++} Begins Low	8.27
Triage+	Schedule	α_+ Begins Low	6.23
Triage+	Schedule	α_+ Begins High	5.56
Triage+	Fixed α_+		4.6
Triage+	Reactive	α_+ Begins High	1.3

Method			Maximum
Triage++	Schedule	α_{++} Begins High	10.94
Triage++	Fixed α_{++}		10.7
Triage++	Reactive	α_{++} Begins High	10.68
Triage++	Schedule	α_{++} Begins Low	10.49
Triage+	Fixed α_+		10.11
Triage+	Reactive	α_+ Begins Low	9.92
Triage+	Schedule	α_+ Begins Low	9.83
Triage+	Schedule	α_+ Begins High	9.74
Triage++	Reactive	α_{++} Begins Low	9.14
Triage+	Reactive	α_+ Begins High	2.38

Next we note that the variable α_{++} Triage++ methods perform best when they begin with a high α_{++} value. This corresponds to beginning with a less restrictive decision posture. If varying α_{++} by schedule the method would then become more restrictive, while the reactive method would adjust the α_{++} as needed. The corresponding posture in the Triage+ methodology would be to begin with a low value of α_+ .

Intuitively one might expect the best performance to come from the reactive methods, and in general that appears to be the case. The overall best average performance comes from the reactive Triage++ method, and the best average performance by a Triage+ method is from the corresponding reactive application. Given that the average portfolio value achieved by the Come and Go method across our simulation runs was 10.3, the 10.21 average achieved by the reactive Triage++ method, which would provide the decision maker added temporal flexibility, is encouraging. This method also achieved a worst portfolio value of 9.78, indicating that it is fairly robust against poor choices of ρ and γ . It cannot be overstated, however, the degree to which this performance relies on a quality CDF estimate.

VII. Conclusions and Suggestions for Further Research

At the outset of this effort, five specific research goals were outlined in Chapter 1 Section 1.2. Over the course of the three original papers presented in this dissertation each has been addressed in turn. Collectively, these form the original contributions of this effort which are applicable to continuous decision problems that arise frequently in a variety of fields including early systems engineering and operations research.

7.1 A Definition of Continuous Decision Problems

Chapter 4 Section 4.1 develops a formal definition for continuous decision problems. In short, it is a decision problem in which the decision maker, within a finite time horizon, expects to sequentially engage in more than one decision event en route to a final selection. Chapter 4 Sections 4.1.1 through 4.1.3 provide the background and motivation for each of the elements of this definition. To date, this appears to be the first attempt to provide a formal definition of this class of problems.

7.2 Effective Methods for Continuous Decision Problems

As originally specified, the Triage Method is designed for screening alternatives and is not suitable for portfolio selection as it does not take resource requirements into account. Two extensions are offered, termed Triage+ and Triage++, to overcome this limitation and extend the triage approach to portfolio selection. The extensions are introduced in Chapter 4 Section 4.3 and further developed throughout. Their effectiveness is demonstrated via application to a small, fictional problem in Chapter 4, and to actual historical data from the AFMC Development Planning effort in Chapters 5 and 6.

7.3 Qualifying Criteria

Criteria which described suitable problems for the application of the Triage extensions were identified. First and foremost is the above definition, whose conditions must be met to ensure that the decision situation is indeed a continuous one. Next, the level of data available to inform the decision making process is considered. As discussed in Chapter 6 Section 6.4.1 the importance of the CDF estimate to the Triage++ methodology argues strongly that this methodology should only be employed when there is a historical basis for the CDF estimate. As was shown, even an experienced decision maker can produce a CDF estimate that results in portfolio selection quality that is very poor.

Next the nature of the decision situation and the expected behavior of the alternatives must be considered. A decision maker may seek to construct a portfolio in a continuous fashion for a variety of reasons, one of the most common being the possibility that alternatives will not remain available should the decision be deferred too long. As discussed in Chapter 5 Section 5.5.1 a general rule of thumb appears to be that the Triage+ extension is viable when the expected departure rate of alternatives is greater than $\frac{2}{3}$ the expected arrival rate. If the Triage++ extension can be used, the fraction appears to be closer to $\frac{1}{2}$.

An important finding discussed in Chapter Section 5.5.3 is that the size of the sensitivity interval employed does not appear to impact the effectiveness of the extensions. This allows that the use of the Triage extension does not color the decision analyst's choices on the size of the sensitivity interval, which should be driven only by the decision maker's level of uncertainty.

7.4 Tradeoff Characterization

The exact nature of the tradeoff that is made when a continuous decision approach is applied is difficult to speak to in general as it is highly specific to the individual decision situation. In general there is obviously the potential for loss in the overall value of the portfolio generated by a continuous approach as the decisions are being made with less

information than in a tradition approach where the entire set of alternatives is known. As discussed in Chapter 5 Section 5.4.3 there is again in temporal flexibility however, as the decision maker is potentially able to commit resources to an alternative as soon as a selection decision is made. This would be a particularly attractive aspect of the continuous approach in a scenario where resources are idled while awaiting selected alternatives. This notion of temporal flexibility is further extended in Chapter 6 Section 6.3.2.3 in which the decision maker is provided an ability to control the target level of resource commitment as a function of time and the methodology attempts to meet the target via its selection decisions.

7.5 Guiding Principles

In addition to the entry criteria referred to above in Section 7.3 elicitation techniques that a decision analyst can use to arrive at cutoff values were investigated in Chapter 6 Section 6.2.2. Though generalizing on the basis of the limited experience described here is not recommended, the decision maker expressed a strong preference for eliciting cutoff values via the specification of a minimally acceptable alternative as opposed to via the use of fictional alternatives. In the same vein, encouraging results were seen in using the exponential and log-normal distributions to approximate CDFs.

Taken together then, this research provides a basic set of guidelines to provide the decision analyst in using the Triage extensions. This list is by no means complete, and there are a number of open research questions still associated with the extensions. It would also be beneficial to replicate the testing conducted to date with additional data sets to verify that the observed performance is inherent in the methodology and not an artifact of a particular data set.

7.6 Future Research

First and foremost the Triage extensions should be investigated with a wider variety of historical data sets to verify their performance. Once these data sets have been identified,

they should be utilized for research on a number of aspects of using the extensions, including:

- Robust methods for eliciting and modeling of CDFs
- Improved methods of setting α values
- Further investigation of the modified calculation of R

Bibliography

- [1] Adkins, R. and D. Paxson. "Evaluating Real Sequential R&D Investment Opportunities." *Real Options Conference*. 2012.
- [2] Ahn, Byeong Seok, et al. "Multi-Attribute Decision Aid Under Incomplete Information and Hierarchical Structure," *European Journal of Operational Research* (2000).
- [3] Aloysius, John, et al. "User Acceptance of Multi-Criteria Decision Support Systems: The Impact of Preference Elicitation Techniques," *European Journal of Operational Research* (2006).
- [4] Anandalingam, G. and C. E. Olsson. "A multi-stage multi-attribute decision model for project selection," *Eur. J. Oper. Res.*, 43:271–283 (1989).
- [5] Anandalingam, G. and C. E. Olsson. "A Multi-Stage Multi-Attribute Decision Model for Project Selection," *European Journal of Operational Research* (1989).
- [6] Asundi, J., et al. *Using Economic Considerations to Choose Among Architecture Design Alternatives*. Technical Report, Carnegie Mellon Software Engineering Institute, 2001.
- [7] Bana e Costa, C. A., et al. "Development of Reusable Bid Evaluation Models for the Portuguese Electric Transmission Company," *Decision Anal.*, 5(1):22–42 (2008).
- [8] Bansal, Saurabh, et al. "Using Expert Assessment to Estimate Probability Distributions," (2014).
- [9] Bard, J., et al. "An Interactive Approach to R&D Project Selection and Termination," *IEEE Transactions on Engineering Management*, 35(3):139–146 (1988).
- [10] Beach, L. R. "Broadening the definition of decision making: The role of prechoice screening of options," *Psychological Sci.*, 4(4):215–220 (1993).
- [11] Benninga, S. and E. Tolokowsky. "Real Options - An Introduction and an Application to R&D Valuation," *The Engineering Economist*, 47(2):151–168 (2002).
- [12] Beyth-Marom, Ruth. "How Probable is Probable? A Numerical Translation of Verbal Probability Expressions," *Journal of Forecasting* (1982).
- [13] Bickel, J. and J. Smith. "Optimal Sequential Exploration: A Binary Learning Model," *Decision Analysis*, 3(1):16–32 (2006).
- [14] Butler, J., et al. "Simulation Techniques for the Sensitivity Analysis of Multi-criteria Decision Models," *European Journal of Operational Research*, 103:531–546 (1997).

- [15] Canbolat, Y., et al. "Combining Decision Tree and MAUT for Selecting a Country for a Global Manufacturing Facility," *Omega*, 35:312–325 (2005).
- [16] Chapanis, A. *Human Factors in Systems Engineering*. Wiley, 1996.
- [17] Chen, Y., et al. "A case-based distance method for screening in multiple-criteria decision aid," *Omega*, 36:373–383 (2008).
- [18] Chen, Ye, et al. "Screening in Multiple Criteria Decision Analysis," *Decision Support Systems* (2008).
- [19] Chien, C. "A Portfolio-Evaluation Framework for Selecting R&D Projects," *R&D Management*, 32(4):359–368 (2002).
- [20] Chinchuluun, A. and P. Pardalos. "A Survey of Recent Developments in Multiobjective Optimization," *Annals of Operations Research*, 154:29–50 (2007).
- [21] Clemen, R. T. and J. E. Smith. "On the choice of baselines in multiattribute portfolio analysis: A cautionary note," *Decision Anal.*, 6(4):256–262 (2009).
- [22] Clemen, Robert and Terence Reilly. *Making Hard Decisions with DecisionTools*. Duxbury, 2001.
- [23] Coderre, M., "JIEDDO Fiscal Year 2011 Budget." Retrieved May, 2011: <https://www.jieddo.dod.mil/article.aspx?ID=854>, 2011.
- [24] Dawley, L. M., et al. *Developing a decision model for Joint Improvised Explosive Device Defeat Organization (JIEDDO) Proposal Selection*. MS thesis, Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, 2008.
- [25] de Neufville, R. *Applied Systems Analysis*. McGraw-Hill Publishing Company, 1990.
- [26] DeTeresa, S. J. and M. W. McErlean. *The Joint Improvised Explosive Device Defeat Organization: DOD's Fight Against IEDs Today and Tomorrow*. U.S. House of Representatives Committee on Armed Services: Subcommittee on Oversight & Investigations, 2008.
- [27] Duncan, K. *An Examination of the Decision Analysis Approach to R&D Portfolios*. MS thesis, Virginia Commonwealth University, 2009.
- [28] Durbach, Ian and Theodor Stewart. "A Comparison of Simplified Value Function Approaches for Treating Uncertainty in Multi-Criteria Decision Analysis," *Omega* (2012).
- [29] Dyer, J. "Multiple Criteria Decision Making, Multiattribute Utility Theory: The Next Ten Years," *Management Science*, 38:645–654 (1992).
- [30] Edwards, W., et al. *Advances in Decision Analysis: From Foundations to Applications*. Cambridge University Press, 2007.

- [31] Edwards, W. and J. R. Newman. *Multiattribute Evaluation*. Beverly Hill, CA: Sage Publications, Inc., 1982.
- [32] Fasolo, Barbara, et al. "Escaping the Tyranny of Choice: When Fewer Attributes Make Choice Easier." London School of Economics Working Paper LSEOR 05.80, 2006.
- [33] Garvey, P. *Analytical Methods for Risk Management: A Systems Engineering Perspective*. CRC Press, 2009.
- [34] Gilbride, T. J. and G. M. Allenby. "A choice model with conjunctive, disjunctive, and compensatory screening rules," *Marketing Sci.*, 23(3):391–406 (2004).
- [35] Godinho, P., et al. "On the use of Multiple Financial Methods in the Evaluation and Selection of Investment Projects," *Investigao Operacional*, 24(1) (2004).
- [36] Guitouni, Adel and Jean-Marc Martel. "Tentative Guidelines to Help Choosing an Appropriate MCDA Method," *European Journal of Operational Research* (1998).
- [37] Gurgur, C. and C. Morley. "Lockheed Martin Space Systems Company Optimizes Infrastructure Project-Portfolio Selection," *Interfaces*, 38(4):251–262 (2008).
- [38] Gutman, A., et al. "The Triage Method: Screening Alternatives Over Time with Multiobjective Decision Analysis." Submitted to International Journal of Multi-Criteria Decision Analysis. Contact jweir@afit.edu for copies., January 2013.
- [39] Haimes, Y. *Risk Modeling, Assessment, and Management*. Wiley, 2004.
- [40] Haimes, Y., et al. "Multiobjective Decision Tree Analysis," *Risk Analysis*, 10:111–127 (1990).
- [41] Hendrix, Jeremy, et al. "Continuous Decision Support," *International Journal of Multicriteria Decision Making* (2014).
- [42] Hendrix, Jeremy, et al. "The Effects of Decision Problem Parameters on the Effectiveness of the Extended Triage Method," *European Journal of Operational Research* (2015).
- [43] Henig, M. and H. Katz. "R&D Project Selection: A Decision Process Approach," *Journal of Multi-Criteria Decision Analysis*, 5:169–177 (1996).
- [44] Henriksen, A. and A. Traynor. "A Practical R&D Project-Selection Scoring Tool," *IEEE Transactions on Engineering Management*, 46(2):158–170 (1999).
- [45] Howard, R. "Decision Analysis: Applied Decision Theory." *Proceedings of the 4th International Conference on Operational Research*. 55–77. 1966.
- [46] Irragori, F López and NM Arriata Martínez. *R&D project portfolio selection in public organizations*. Technical Report, Tech. Rep. PISIS-2013-01, Graduate Program in Systems Engineering, UANL, San Nicolás de los Garza, México, 2013.

- [47] JIEDDO, “About JIEDDO.” Retrieved July, 2009: <https://www.jieddo.dod.mil/aboutJIEDDO/ajhome.aspx>, 2009.
- [48] Johnson, P., et al. *Enterprise Architecture Analysis with Extended Influence Diagrams*. Royal Institute of Technology , Stockholm, Sweden, 2007.
- [49] Joint Improvised Explosive Device Defeat Organization. *Annual Report*. Technical Report, JIEDDO, 2010.
- [50] Kahraman, Y. *Robust Sensitivity Analysis for Multi-Attribute Deterministic Hierarchical Value Models*. MS thesis, Air Force Institute of Technology, 2002.
- [51] Keefer, Donald and Samule Bodily. “Three-Point Approximations for Continuous Random Variables,” *Management Science* (1983).
- [52] Keeney, R. and H. Raiffa. *Decisions With Multiple Objectives: Preferences and Value Tradeoffs*. Wiley Publishing, 1976.
- [53] Keeney, R. L. *Siting Energy Facilities*. New York, NY: Academic Press, 1980.
- [54] Keeney, R. L. *Value-Focused Thinking: A Path to Creative Decisionmaking*. Cambridge, MA: Harvard University Press, 1992.
- [55] Keeney, R. L. “Common mistakes in making value trade-offs,” *Oper. Res.*, 50(6):953–945 (2002).
- [56] Keeney, Ralph L and Detlof von Winterfeldt. *Practical value models*, chapter 13, 232–252. Cambridge University Press, 2007.
- [57] Keisler, J. “Value of Information in Portfolio Decision Analysis,” *Decision Analysis*, 1(3):177–189 (2004).
- [58] Kira, D., et al. “A Specific Decision Support System (SDSS) to Develop an Optimal Project Portfolio Mix Under Uncertainty,” *IEEE Transactions on Engineering Management*, 7(3):213–221 (1990).
- [59] Kirkwood, C. W. *Strategic Decision Making: Multi-Objective Decision Analysis with Spreadsheets*. Belmont, CA: Duxbury Press, 1997.
- [60] Köksalan, M., et al. “A flexible approach to ranking with an application to MBA programs,” *Eur. J. Oper. Res.*, 201:470–476 (2010).
- [61] Kozat, S. and A. Singer. “Universal Portfolios with Switching and Side Information.” Submitted to IEEE International Conference on Acoustics, Speech, and Signal Processing, May 2006.
- [62] Liesio, J. *Robust Portfolio Optimization in Multi-Criteria Project Selection*. MS thesis, Helsinki University of Technology, 2006.

- [63] Liesiö, J., et al. "Preference programming for robust portfolio modeling and project selection," *Eur. J. Oper. Res.*, 181:1488–1505 (2007).
- [64] Lim, C., et al. "Sequential Search with Multiattribute Options," *Decision Analysis*, 3(1):3–15 (2006).
- [65] Limpert, Eckhard, et al. "Log-normal Distributions Across the Sciences: Keys and Clues," *BioScience* (2001).
- [66] Lootsma, F. *Multi-Criteria Decision Analysis via Ratio and Difference Judgment*. Kluwer Academic Publishers, 1999.
- [67] Malakooti, B. "Ranking and screening multiple criteria alternatives with partial information and use of ordinal and cardinal strength of preferences," *IEEE Trans. on Systems, Man, and Cybernetics Part A*, 30(3):355–368 (2000).
- [68] Marshall, Kneale and Robert Oliver. *Decision Making and Forecasting*. McGraw-Hill, 1995.
- [69] Mild, P. and A. Salo. "Combining a Multiattribute Value Function with an Optimization Model: An Application to Dynamic Resource Allocation for Infrastructure Maintenance," *Decision Analysis*, 6(3):139–152 (2009).
- [70] Miller, L. and C. Park. "Decision Making Under Uncertainty - Real Options to the Rescue?," *The Engineering Economist*, 47(2):105–150 (2002).
- [71] Montgomery, Douglas, et al. *Introduction to Time Series Analysis and Forecasting*. Wiley, 2008.
- [72] Morton, Alec, et al. "Multicriteria Portfolio Decision Analysis for Project Selection." London School of Economics Working Paper LSEOR 13.141, 2013.
- [73] Mustajoki, J., et al. "Using intervals for global sensitivity and worst-case analyses in multiattribute value trees," *Eur. J. Oper. Res.*, 174:278–292 (2006).
- [74] Parnell, G., et al. *Decision Making in Systems Engineering and Management*. Wiley, 2008.
- [75] Phillips, Lawrence D. and Carlos A. Bana e Costa. "Transparent prioritisation, budgeting and resource allocation with multi-criteria decision analysis and decision conferencing," *Annals of Operations Research*, 154(1):51–68 (2007).
- [76] Polyashuk, M. "A Formulation of Portfolio Selection Problem with Multiple Criteria," *Journal of Multi-Criteria Decision Analysis*, 13:135–145 (2005).
- [77] Sarin, R. K. "Screening of multiattribute alternatives," *Omega*, 5(4):481–489 (1977).
- [78] Sevastjanov, P. and L. Dymova. "Stock screening with use of multiple criteria decision making and optimization," *Omega*, 37:659–671 (2009).

- [79] Skinner, D. *Decision Analysis: A Practitioner's Guide to Improving Decision Quality*. Probabilistic Publishing, 2009.
- [80] Smith, J. and R. Nau. "Valuing Risky Projectys: Option Pricing Theory and Decision Analysis," *Management Science*, 41(5):795–816 (1995).
- [81] Spradlin, C. T. and D. M. Kutoloski. "Action-oriented portfolio management," *Res. Tech. Management*, 42(2):26–32 (1999).
- [82] Spradlin, C. Thomas and David Kutoloski. "Action-Oriented Portfolio Management," *Research-Technology Management* (1999).
- [83] Stummer, C. and K. Heidenberger. "Interactive R&D portfolio analysis with project interdependencies and time profiles of multiple objectives," *IEEE Trans. on Engineering Management*, 50(2):175–183 (2003).
- [84] Talias, M. "Optimal Decision Indices for R&D Project Evaluation in the Pharmaceutical Industry: Pearson Index versus Gittins Index," *European Journal of Operational Research*, 177:1105–1112 (2007).
- [85] Walker, W. E. "The use of screening in policy analysis," *Management Sci.*, 32(4):389–402 (1986).
- [86] Weir, Jeffery, et al. "The Triage Method: Screening Alternatives Over Time with Multiobjective Decision Analysis," *International Journal of Multicriteria Decision Making* (2014).
- [87] West, P. D. *Decision Making in Systems Engineering and Management*, chapter Solution Design, 241–282. Hoboken, NJ: Wiley & Sons Inc., 2011.
- [88] Wilson, C. *Improvised Explosive Devices (IEDs) in Iraq and Afghanistan: Effects and Countermeasures*. Technical Report, Congressional Research Service, 2007.
- [89] Yoon, K. P. and C. -L. Hwang. *Multiple Attribute Decision Making: An Introduction*. Thousand Oaks, CA: Sage Publications, 1995.
- [90] Zaroliagis, Christos. "Recent Advances in Multiobjective Optimization." *Stochastic Algorithms: Foundations and Applications 3777*. Lecture Notes in Computer Science, edited by OlegB. Lupanov, et al., 45–47, Springer Berlin Heidelberg, 2005.

REPORT DOCUMENTATION PAGE					<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY) 24-12-2015		2. REPORT TYPE Dissertation		3. DATES COVERED (From — To) October 2008 - December 2015		
4. TITLE AND SUBTITLE Continuous Decision Support				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Hendrix, Jeremy P., GS-14, DAFC				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way Wright-Patterson AFB, OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENV-DS-15-D-018		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Intentionally Left Blank				10. SPONSOR/MONITOR'S ACRONYM(S)		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED						
13. SUPPLEMENTARY NOTES This work is declared a work of the U.S. Government and is not subject to copyright protection in the United States.						
14. ABSTRACT Organizations are often faced with portfolio construction efforts that require them to select one or more alternatives, subject to resource constraints, with the aim of achieving the maximum value possible. This is a well-defined problem with a number of analytically defensible approaches, provided the entire set of alternatives is known when the decision event takes place. Less well treated in the literature is how to approach this problem when the entire set of alternatives is unknown, as when the alternatives arrive over time. This change in the availability of data shifts the problem from one of identifying an optimal subset to one in which a series of smaller decisions are undertaken regarding the acceptability of each alternative as it presents itself. This work expands upon a methodology known as the Triage Method. The original Triage Method provided a screening tool that could be applied to alternatives as they presented themselves to determine if they should be accepted for further study, rejected out of hand, or held pending until later date. This decision was made strictly upon the value of the alternative and with no consideration of its cost. Two extensions to the Triage Method are offered which provide a capability to consider cost and other resource requirements of the alternatives, thus allowing a move from simply screening to portfolio selection. Guidelines are presented as to when each of these extensions is best employed, a characterization of the performance tradeoff between these and more traditional methodologies is developed, and insight and techniques for setting the value of parameters required by the extensions are provided.						
15. SUBJECT TERMS Decision analysis; portfolio; cost/benefit analysis						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT		18. NUMBER OF PAGES	
a. REPORT	b. ABSTRACT	c. THIS PAGE			19a. NAME OF RESPONSIBLE PERSON Dr. David R. Jacques (ENV)	
U	U	U	UU		157	
						19b. TELEPHONE NUMBER (include area code) (937) 255-3355 x3329 David.Jacques@afit.edu