

Introduction to PALMA™: A Tool for Portfolio Investment Analysis

September 2020

MITRE Product (MP) 200652

PALMA


A Tool for Portfolio Investment Analysis

Introduction

MITRE's Portfolio Analysis Machine (PALMA™)¹ is a graphical and computational framework to help decision-makers select the best portfolio of investments, based on cost and operational benefit, so they can best leverage limited resources. It has been in use since 2000 to conduct tradeoff and portfolio optimization analyses for multiple government organizations. Over that time, it has continually evolved to meet sponsor needs by incorporating new features, functions, and computational capabilities.

PALMA Overview

PALMA is designed to solve the “knapsack problem,” as illustrated in the figure below: Given the constraint of limited space in the knapsack, what combination of objects (i.e., portfolio) will provide a camper with the greatest utility? Campers would need to consider dependencies among items (e.g., a can opener is required if the camper intends to bring food in cans) as well as potential substitutes (e.g., candles versus a battery-powered lantern) in their decision-making process.



¹ PALMA™ is a trademark of The MITRE Corporation.

PALMA is built on the principle that the process of making investment decisions should begin by asking basic questions such as:

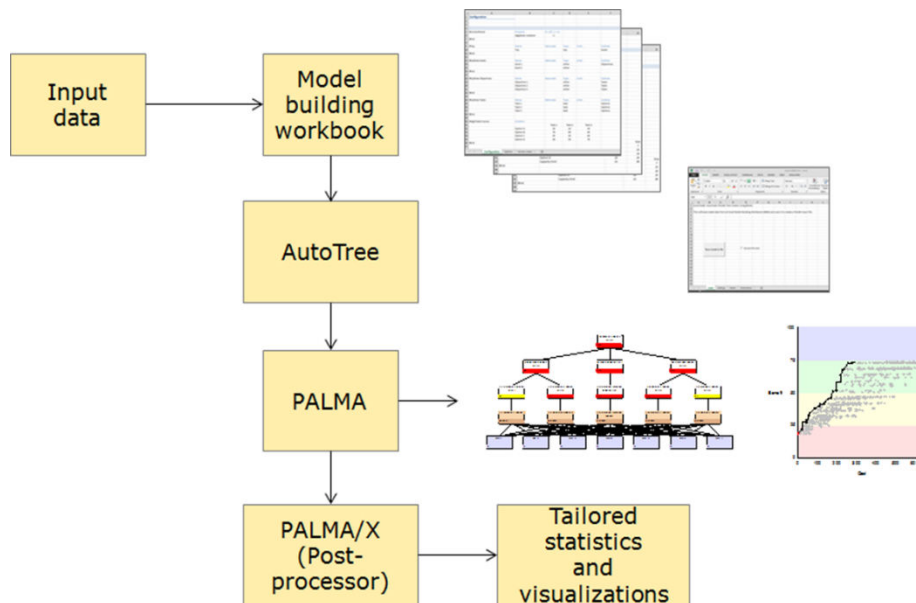
- What is the decision to be made (e.g., what is the need to be met?)
- How well is the need being met today? What gaps or shortfalls exist?
- What investment options should be considered to fill the gaps or remediate the shortfalls?
- How much do they cost and how much benefit do they provide?
- How do these investments contribute to the organization’s overall goal?

PALMA provides (1) a language in which the answers to these questions can be expressed and (2) a framework for combining information to gain insight into the investment problem. This framework provides a structured, repeatable approach to investment planning.

PALMA’s full toolset, typical uses, and fundamental concepts will be discussed in detail below. A notional search and rescue mission will be used to illustrate the key concepts.

The PALMA Toolset

The main element of the PALMA toolset is the PALMA application itself, but the toolset also includes pre- and post-processing tools to help in building the model and analyzing the results, as illustrated in the figure below. The model building workbook (MBW) and AutoTree are pre-processing tools that allow users to build the model in Microsoft Excel® using the input data they have developed. The post-processor tool, PALMA-X, allows users to automate repetitive tasks and perform post-processing of model results, including sensitivity analysis and graphical display. The pre- and post-processing elements of the toolset are discussed further in later sections of this paper.



In addition, the toolset includes comprehensive user’s manuals for PALMA and AutoTree as well as a large set of example files.

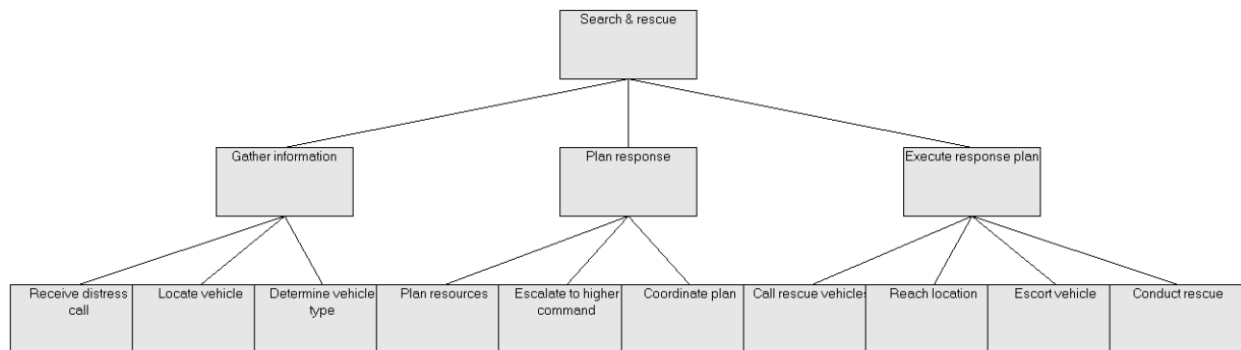
Fundamental PALMA Model Concepts

PALMA is based on the principles of multi-attribute utility theory (MAUT), a structured decision-making methodology for guiding tradeoffs among multiple, sometimes competing, objectives (Reference 1). The investment options are assessed and compared based on how well they satisfy the defined objectives, not on financial metrics like cost savings, net present value, or return on investment. In a PALMA model, cost and benefit are considered separately; that is, cost is an independent variable and is not included as an objective in the benefit calculation. This allows PALMA to build a benefit/cost graph that illustrates the benefit across a range of costs, which is useful in making trade-offs.

The Value Tree

A PALMA model is based on the concept of a value tree. The value tree identifies the high- level mission and decomposes it into goals, objectives, actions, and tasks (or another hierarchical or functional breakdown). Each of the elements in the value tree is called a node. The value tree hierarchy also shows the relationships among the nodes (also known as parent-child relationships). The value tree is decomposed to the level at which the effect of the investment options can be directly evaluated (i.e., leaf nodes), with the top node representing the overall mission.

A simple value tree for a notional search and rescue mission is shown in the figure below. In this example, the overall mission of search and rescue is broken down into three key tasks: gather information, plan response, and execute response plan. Each task has multiple subtasks; for example, gather information depends on the receive distress call, locate vehicle, and determine vehicle type subtasks.



A value tree can be built by subject matter experts and/or by leveraging existing organizational resources reflecting its objectives or goals (e.g., an operational architecture, a business reference model, a strategic plan, mission threads, etc.).

Scoring and Utility

Every node in the PALMA value tree has a numerical score. The score of a leaf node represents the impact of the selected portfolio of investment options on that leaf. The score of a non-leaf node represents the rollup (combination) of the scores of all that node’s children (see section below for more about rollup rules). The score at the top of the tree is the overall benefit score for the selected portfolio; this benefit score, along with the portfolio cost, is what PALMA uses in its optimization calculations.

The node scores, in general, represent utility, a dimensionless value on a 0-100 scale where zero represents no capability and 100 represents the ideal capability. The exception is engineering nodes, whose scores represent physical units (e.g., vehicle speed in miles per hour, time delay in delivering a data message). These scores must be converted into utility scores before they can be combined with the scores of other nodes.

PALMA associates colors with the numerical scores to provide a quick visual indication of node status when viewing the value tree. PALMA comes with a standard set of color files, and the user can also define custom color schemes. As an example, the color scheme for the search and rescue model example used in this paper is shown below.

Color	Description	Score Range
Blue	Exceeds requirements	$score \geq 90$
Green	Fully meets requirements	$75 \leq score < 90$
Yellow	Partially meets requirements	$50 \leq score < 75$
Orange	Partially meets requirements, but with significant gaps	$25 \leq score < 50$
Red	Meets very few requirements	$1 \leq score < 25$
Zero	Does not meet requirements	$0 \leq score < 1$

Rollup Functions

Rollup functions mathematically characterize the relationships between connected nodes. These functions specify how the scores of the “child” nodes combine to produce the score of the “parent” node. The complete set of rules, at all levels of the value tree, allows the benefit at the top of the tree to be calculated from the scores of the leaf nodes.

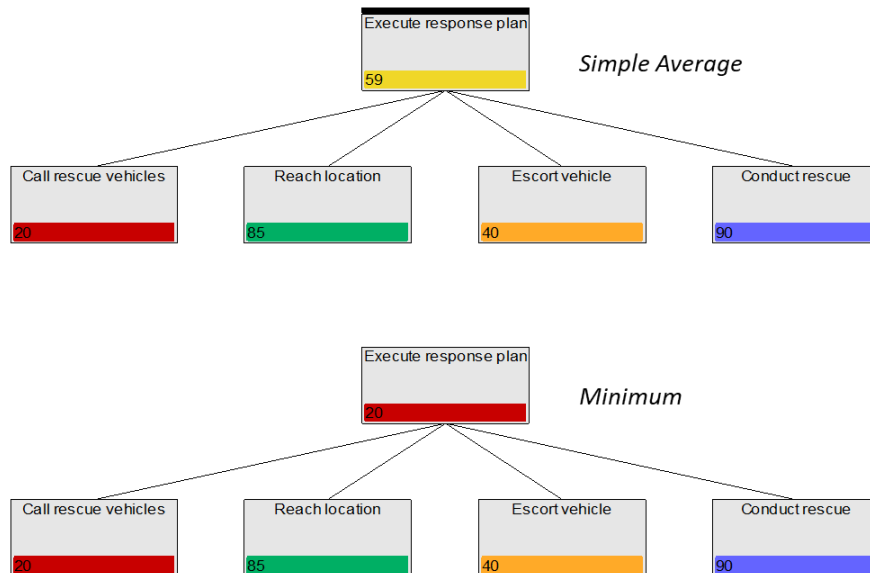
PALMA offers a wide variety of built-in roll-up functions including:

- Maximum, minimum, arithmetic average, weighted average, advanced averages
- Curve fitting techniques
- Lookup tables

PALMA also provides flexibility to customize rollup functions using algebraic notation and standard math functions. Rollup functions are typically determined through

discussion with stakeholders who have a deep understanding of the overall goal/objective/mission.

The figure below shows the results of using two different rollup functions on the same set of scores.



Baseline Scores

Baseline scores reflect the current state of your mission or overall goal without any of the defined investment options being funded. Baseline scores are assigned to each leaf node and are propagated up the value tree based on the defined rollup rules. From this, you can see where you are doing well and where there are shortfalls and/or gaps that may need to be addressed. You can examine how the baseline drives value at the top of the tree.

In some cases, it may not be possible to develop baseline scores, for a variety of reasons, such as having a new mission under evaluation or because the information needed to assign baseline scores is not available. In these cases, the baseline scores of all leaf nodes can be set to zero (which is the default in PALMA).

Investment Options


The investment options represent the alternatives that will be explored in the analysis to support the investment decision being made.

Option Costs. The option costs are determined outside of PALMA and are an input to the model. It doesn't matter what type of costs are used (e.g., acquisition cost or lifecycle cost) or what units they are in; however, for accurate results it is important that all costs be of the same type and in the same units. Costs can include people,

processes, policies, technologies, equipment, materials, facilities, and/or data. In certain types of models, an option’s cost may be set to zero if no resources are needed to realize its impact.

Option Impacts. The option impacts are assessed at the lowest level of the value tree (the leaf nodes). Qualitative or quantitative metrics, or a mix of both, can be used to make these assessments. For each metric, a scoring scheme must be established to define how the impact will be assessed.

For metrics that are measurable (e.g., vehicle speed in miles per hour), a utility function must be defined to translate these values into dimensionless benefit scores. This allows scores for different metrics (e.g., speed in miles per hour and fuel efficiency in miles per gallon) to be combined. It also establishes “how much is good enough”: twice as much isn’t necessarily twice as good, and it may cost a lot more.



A metric that is not measurable can be assessed using a constructed scale. Constructed scales are verbal, descriptive scales used to assess impacts in cases where natural scales do not exist (Reference 2). This type of scale establishes a consistent way of assessing different levels of capability for something that is not directly quantifiable.

For a PALMA model, a numerical utility value must be assigned to each level of the constructed scale, as in the example below.

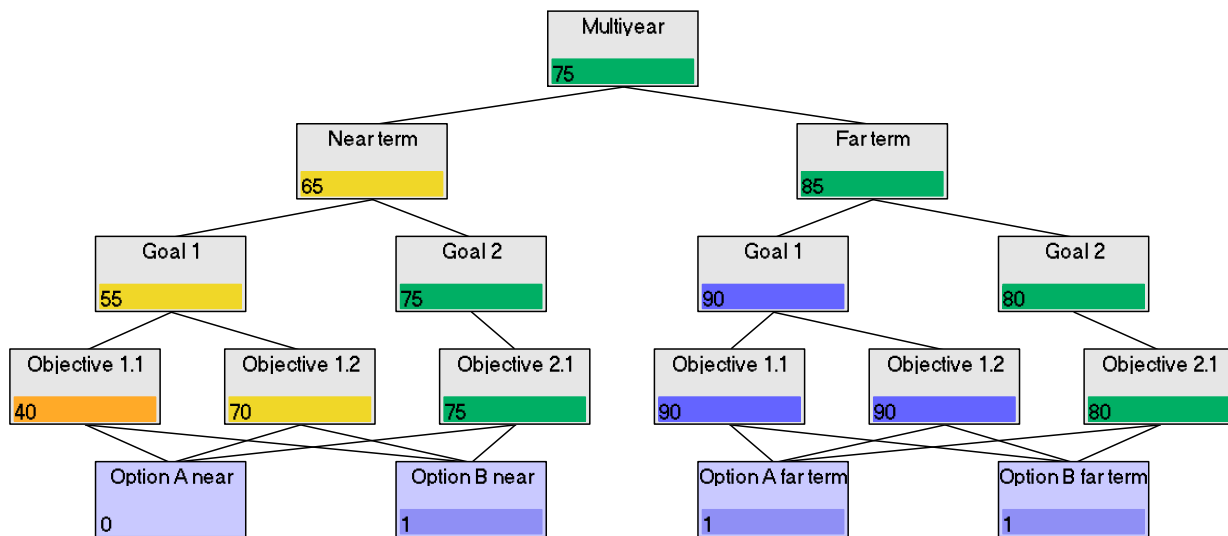
Score	Title	Description
100	Fully satisfied	Meets all requirements
80	Good	Meets all major requirements, with minor limitations
60	Fair	Meets most major requirements, with significant limitations
40	Poor	Fails to meet many major requirements, but provides some value
20	Very poor	Fails to meet many major requirements, but provides minor value
0	No capability	Provides no value

Option Dependencies. Option dependencies are used to capture relationships among the options. Dependencies ensure that the portfolios considered by PALMA take real-world constraints into account. For example, if Option B is an add-on that enhances the

performance of Option A, but is of no value by itself, a dependency should be created to tell PALMA not to include Option B in a portfolio unless Option A is also in that portfolio. PALMA includes several types of dependencies to allow for modeling different types of relationships among investment options.

Multi-Period Modeling

A multi-year (or multi-period) model can be developed in PALMA by creating a separate value tree for each time period being modeled (e.g., near-term and far-term). These separate trees are then connected to a single “multiyear” parent node at the top of the tree, as shown in the figure below. The multi-period structure makes it possible for an option to have different impacts and different costs in different time periods. Usually the basic value tree structure is the same for all time periods, but PALMA can also accommodate unique value trees in each time period if needed.



Creating a PALMA Model

All the elements of a PALMA model (value tree structure, rollup functions, investment option costs and impact, etc.) are stored in a text file called a TRE file. The TRE file can be built using either PALMA’s graphical user interface (GUI) or the AutoTree pre-processing tool.

In the PALMA GUI, buttons and text boxes allow the user to

- Define the tree structure (add, delete, and move nodes; define node rollup rules and other properties; change the parent of a child node; copy and “paste” a node or subtree);
- Define the investment options (cost, impact, dependencies).

Using the GUI can be a great way to start becoming familiar with PALMA. It can become cumbersome, however, for large or complex models.

The AutoTree pre-processor is a macro-enabled Excel workbook that can be used to automate construction of the PALMA input file. To use AutoTree, the data elements that make up the model are entered in an Excel file called a model-building workbook. The MBW consists of a series of formatted blocks representing the various model elements (e.g., the top node in the tree, the child nodes of a given parent node, and the investment options). AutoTree then reads the information from the MBW and uses it to generate a properly formatted PALMA input file.

Using the PALMA Model and Analyzing Results

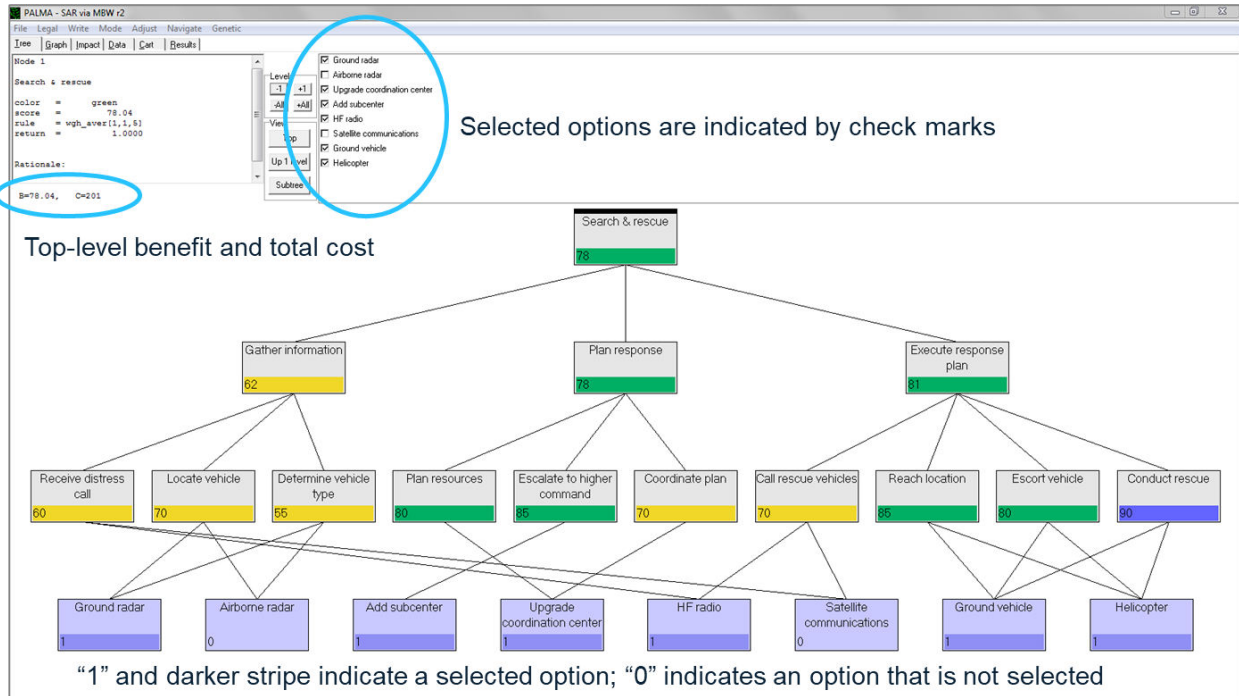
Once the model has been created, PALMA can be used to investigate individual investment options or portfolios or to run optimization algorithms that identify portfolios that provide the greatest benefit as a function of cost.

The user can choose one of three optimization methods: exhaustive search of all allowable combinations of options (if the number of options is less than 30) and two genetic algorithms, iterative mathematical algorithms that mimic the processes of biological evolution. Based on the total cost and top-level of each portfolio examined, PALMA uses the selected optimization algorithm to derive the efficient frontier, the set of portfolios that provide the greatest benefit for a given cost.

PALMA's main display options are the "Tree" page, which displays the value tree hierarchy and node scores and the "Graph" page, which displays a benefit-cost graph of the efficient frontier and a representative set of less cost-effective portfolios.

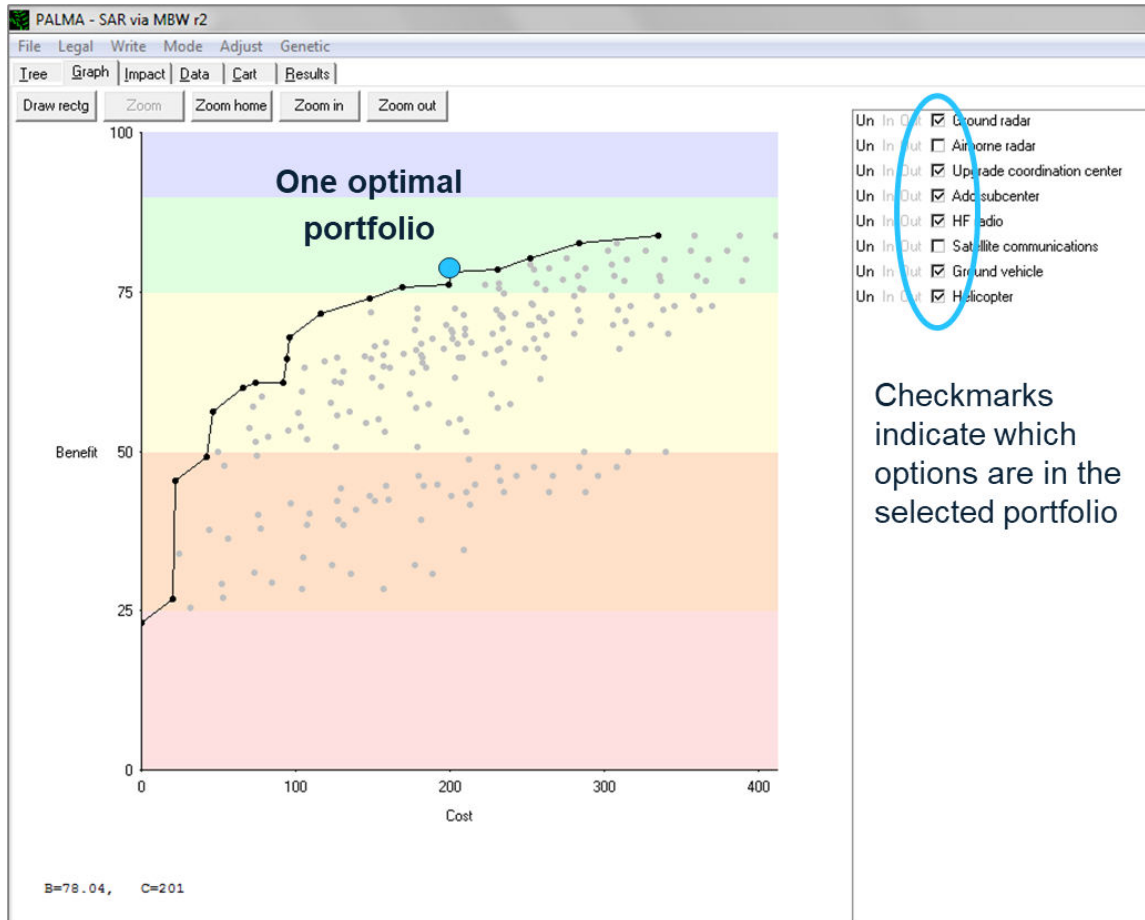
Exploring the Value Tree (Tree Page)

The "Tree" page can be used to investigate individual options or portfolios, as illustrated below, by checking the appropriate boxes in the list of options in the upper right-hand field of the display. The impact of selecting these options and the resulting rolled-up benefit at each node in the tree, can easily be seen by the node scores and the color bars. This allows the user to identify where requirements are being met and where there may be gaps.



Exploring the Efficient Frontier (Graph Page)

The "Graph" page can be used to explore the efficient frontier (black dots in the figure below), which identifies portfolios (and the investments in each) that provide the greatest benefit at a specific funding level. The graph also illustrates a representative set of less efficient portfolios (gray dots in the figure below.)



Other Analysis Techniques

PALMA results (portfolios, node scores, efficient frontiers) can also be exported to a tab-delimited text file and converted to Excel for additional analysis. An example of the type of analysis that can be done is to examine how often each investment appears on the efficient frontier, considered across all frontier points or a user-defined cost range.

The PALMA toolset also includes automation and post-processing capability using PALMA/X, a dynamic link library (DLL) that allows applications such as Excel or MATLAB to use macros to call PALMA functions (e.g., create the efficient frontier). PALMA/X allows users to automate repetitive tasks in PALMA and perform post-processing of model results, e.g., sensitivity analysis. It can also be useful for integrating PALMA with other analysis tools.

In addition, both PALMA and AutoTree employ extensive error-checking when reading input files and provide detailed error messages to help the user in troubleshooting.

Other Uses of PALMA

While PALMA was designed to support portfolio analysis, its inherent capability to conduct tradeoffs amongst multiple options will support analyses where the objective is

to pick one of several options including analysis of alternatives and cost-benefit, cost capability, and business case analyses. PALMA can also support risk analysis where options are mitigations and benefit is the “burn down” of risk. PALMA’s graphic displays are also useful in the visualization of complex decision spaces.

Lessons Learned

Lessons learned from twenty years of experience using PALMA include the following:

- Let everyone know up front that PALMA depends largely on expert judgment, though expert judgment may be supplemented with modeling and simulation analyses and other authoritative information in some studies.
- Make sure the decision that the model is intended to support is as well defined as possible.
- If possible, tie in your analysis with existing processes in your organization.
- Try to get buy-in from decision-makers and stakeholders as early as possible.
 - Start with a small “pilot” effort and build from there, using the results from the pilot and the feedback from the stakeholders to refine your approach as you go along.
 - It often helps to develop an initial, notional model leveraging any available artifacts to start the process. It is usually easier for stakeholders to react to a strawman model than to create a model out of whole cloth.

Summary

PALMA is a flexible and powerful tool that can be used to model a wide range of investment problems. Its optimization algorithms automatically identify investment portfolios that provide the most “bang for the buck,” even for large numbers of investment options. A PALMA model provides a framework for linking investment options to improved capability or mission effectiveness; this gives insight into the impact of investment choices and presents the big picture to the decision-maker.

Perhaps most important, using PALMA establishes a structured, repeatable, transparent process for approaching an investment decision and thus provides a basis for justifying the resulting decision.

Inquiries about PALMA should be sent to palma@mitre.org

References

- [1] Keeney, R.L., and Raiffa, H., 1976, *Decisions with Multiple Objectives: Preference and Value Tradeoffs*, John Wiley, New York, NY.
- [2] Garvey, P.R., October 2008, *Analytical Methods for Risk Management: A Systems Engineering Perspective*, CRC Press, New York.



About MITRE

MITRE's mission-driven teams are dedicated to solving problems for a safer world. Through our public-private partnerships and federally funded R&D centers, we work across government and in partnership with industry to tackle challenges to the safety, stability, and well-being of our nation.