



ACQUISITION RESEARCH PROGRAM SPONSORED REPORT SERIES

A Robust Framework for Analyzing Acquisition Alternatives

15 October 2018

Moshe Kress, Distinguished Professor

Brian Morgan, CAPT USN

Operations Research Department

Naval Postgraduate School

Approved for public release; distribution is unlimited.

Prepared for the Naval Postgraduate School, Monterey, CA 93943.



The research presented in this report was supported by the Acquisition Research Program of the Graduate School of Business & Public Policy at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact any of the staff listed on the Acquisition Research Program website (www.acquisitionresearch.net).



ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

Abstract

Analysis of Alternatives (AoA) is a complex, multi-dimensional analytical comparison of the operational effectiveness, suitability, and life cycle cost of material solutions to satisfy established capability needs. It is a study completed in the capability development and acquisition processes within the U.S. Department of Defense acquisition system. The four dimensions of the procedure are alternatives, criteria, scenarios, and stakeholders. Current AoA studies lack the structural rigor needed from such a complex procedure, which involves measurements, evaluations, analyses, and modeling, as well as social aspects—all in a highly uncertain environment. We propose a structured paradigm for conducting AoA, rooted in well-established methods of multi-criteria decision analysis. The core of the methodology comprises the concepts of ratio-scale matrices and the Euclidean norm. The ratio-scale matrices are used to elicit evaluations, preferences, and opinions from individual stakeholders and analysts, and the Euclidean norm is utilized to remove preference inconsistencies and help form consensus.



THIS PAGE LEFT INTENTIONALLY BLANK





ACQUISITION RESEARCH PROGRAM SPONSORED REPORT SERIES

A Robust Framework for Analyzing Acquisition Alternatives

15 October 2018

Moshe Kress, Distinguished Professor

Brian Morgan, CAPT USN

Operations Research Department

Naval Postgraduate School

Disclaimer: The views represented in this report are those of the authors and do not reflect the official policy position of the Navy, the Department of Defense, or the federal government.



THIS PAGE LEFT INTENTIONALLY BLANK



Table of Contents

1.	Introduction.....	1
2.	The Four Main Components of Acquisition AoA	5
2.1	Alternatives	5
2.2	Criteria	6
2.3	Scenarios.....	11
2.4	Stakeholders	12
3.	The Set of Criteria.....	13
3.1	Criteria Tree	13
3.2	Criteria Properties	14
3.3	Weights of Criteria	15
4.	Methodology	17
4.1	Determining Criteria Weights for a Certain Scenario	17
4.2	Determining Alternatives' Values for a Certain Criterion and Scenario	19
4.3	The Alternative's Value Function in a Scenario.....	20
4.4	Scenarios' Probabilities.....	20
5.	Selecting the Winning Alternative	23
6.	Implementation	27
	References.....	29



THIS PAGE LEFT INTENTIONALLY BLANK



1. Introduction

The U.S. Department of Defense (DoD) acquisition decision support system comprises three interconnected processes that start with the Joint Capabilities Integration and Development System (JCIDS). Here, shortfalls or gaps in operational capabilities are identified and possible material and non-material approaches are suggested. The second stage, called the defense acquisition system, focuses on determining appropriate materiel solutions for the requirements specified in the first stage. The third stage allocates resources to the decisions made in the second stage. It includes funding and control activities contained in the Planning, Programming, and Budgeting Execution (PPBE) process (DoD, 2017).

The first stage comprises tactical and operational analyses based on wargames, simulations combat models, and input from subject matter experts. It identifies gaps in current capabilities and projects future needs based on evolving threats, operational postures, and changes in strategy. The third stage is a process that allows the Department of Defense to allocate its resources.

The second stage—the defense acquisition system—contains two interrelated phases: a *creative* phase and an *analytic* phase. The outcome of the creative phase is a set of potential materiel solutions to the operational requirements specified in the Initial Capabilities Document (ICD) produced in the first stage. This set comprises acquisition alternatives to be analyzed in the analytic phase. The goal is to consider a comprehensive set of alternatives representing all reasonable solutions. In other words, each alternative in the set must contribute to closing the capability gap(s) encompassing the representative tactical situations while remaining within the bounds of feasibility as specified in the ICD. The analytic phase focuses on evaluating the alternatives with respect to several criteria while incorporating quantitative analysis with multiple stakeholders' opinions and preferences. The outcome of this phase is a recommendation on the most preferred alternative to be considered for acquisition. Such a “winning” alternative may not be unique. This recommendation must be based on multi-criteria



evaluations of the alternatives and reflect a consensus among stakeholders' opinions and preferences.

This study focuses on the analytic phase of the second stage—also called Analysis of Alternatives (AoA; RAND Project Airforce, 2006). Our objective is to propose a comprehensive formal framework and analytic structure for executing AoA. The proposed framework is general enough to be easily tuned to any specific AoA study in any branch of the armed services.

The core of the analytic phase is a Multi-Criteria Decision Analysis (MCDA) process. In this process, alternatives are evaluated according to a set of criteria and the resulting evaluations are then aggregated into a rating or a score that represents the relative standing of each alternative. In a DoD acquisition context, the criteria typically include scenario-dependent operational effectiveness, technological feasibility and risk, supportability, compatibility (with existing systems), and cost. While the general spirit of MCDA is indeed present in typical current AoA projects, its actual manifestation varies significantly among studies (RAND Project Air Force, 2006; TRADOC Analysis Center, 2011; Department of the Navy, 2007; Souders, Schulze, Ginburg, & Goetke, 2004; Smith, Ritz, Smith, & Braun, 2011; RAND, 2016). Crucially, most of these studies lack the structural or formal rigor that is desired in a critical decision process such as AoA. Typical subjects that are either not addressed or treated inconsistently are: systematically evaluating alternatives with respect to criteria and determining the weights of criteria, explicitly treating uncertainty and risk, and adequately aggregating preferences among stakeholders. Moreover, as much as it is an analytic process, AoA is also a social process that involves several (sometimes many) stakeholders. Different stakeholders, representing various parts of the DoD, may have different opinions or preferences regarding the importance (i.e., weights) of criteria. They may differ in their assessments about the likelihood of future scenarios and they may also disagree about the values of alternatives with respect to qualitative (subjective) criteria, where measures of performance (MOPs) and/or measures of effectiveness (MOEs) are either difficult to compute or do not exist altogether. Even measurable (objective) criteria, such as detection range, velocity and firing accuracy (say, probability of hit), may be scaled differently by different stakeholders. This important social aspect seems to be ignored in



current AoA studies. In our proposed framework we attempt to remedy these, and other, shortcomings.

The main contributions of this paper are: (a) proposing a clear “standard” for conducting acquisition AoA in the US Department of Defense, (b) explicitly addressing the role of scenarios and stakeholders in the AoA process, and (c) developing an all-inclusive distance-based model that addresses, simultaneously, all four dimensions of AoA.

The paper is organized as follows. In Section 2 we review the four main ingredients that comprise an acquisition AoA: *alternatives*, *criteria*, *scenarios*, and *stakeholders*. Section 3 describes the set of criteria relevant to a typical DoD acquisition AoA, and its imbedded hierarchy. Section 4 is the heart of this paper, in which we formally describe the MCDA methodology underlying the AoA process. Section 5 addresses the uncertainty associated with future scenarios and the way it affects the choice of the winning alternative. Finally, Section 6 outlines the six steps of a robust AoA study.

In the rest of the paper, we refer to the subject of the AoA study as *item*. An item can be a Navy fighter, a radar system, a transport vehicle, a supply ship, a command, or control system, etc. The objective of the AoA process is to select for acquisition the most appropriate item out of a set of alternative items.



THIS PAGE LEFT INTENTIONALLY BLANK



2. The Four Main Components of Acquisition AoA

Once the operational needs and/or capability gaps have been identified within the JCIDS process, an initial set of items (potential materiel or system solutions to these needs) is generated at the onset of the next process — the defense acquisition system. Generally speaking, the ultimate goal of the acquisition process is selecting, out of an initial set of possible items, an item that provides the best balance between (in-context) utility or value and potential cost and risk. The members of this set of items are called *alternatives*. The alternatives represent the first main component of the acquisition AoA. The other three main components are *criteria* (the touchstones according to which alternatives are evaluated), *scenarios* that provide the operational backdrop for the evaluation, and *stakeholders* who contribute analytic inputs, as well as preferences, opinions, and judgments, into the acquisition decision process.

2.1 Alternatives

Generating the initial set of alternatives – a process called *Development Planning* – is the creative part in acquisition AoA. This process encompasses engineering analysis and technical planning activities that provide the foundation for informed investment decisions. Proposed materiel solutions may represent various stages of maturity in the research-and-development stage, from a completely new system to an existing one.

In some AoA studies there exists a legacy item (ship, weapon, C2 system, etc.) that is either near its end of life, or its capabilities are insufficient for emerging requirements. In such cases, it is important to clearly identify the characteristics of the legacy alternative, which can be considered as a baseline according to which potential upgrades are considered [8].

The set of alternatives should be carefully constructed. It must be not trivial (e.g., just two alternatives where one clearly dominates the other), but also manageable in size. There is hardly an effective and meaningful way of handling the evaluation of dozens of alternatives. One way of reducing the size of the alternatives' set is eliminating similar



alternatives (alternatives that differ marginally or those that are evidently dominated by other alternatives).

The alternatives should also be realistic in the sense that they are technologically feasible and grounded in industry's ability to mature the technology in an acceptable time period. The set of alternatives should not include idealized items that have no practical military value. The set of acquisition alternatives may be divided into categories:

- Modified items currently in operation,
- As-is, or modified, off-the-shelf, items available in the market but not yet in operational use,
- Repurposing and/or recombining existing items with new technologies, and
- Newly developed items. [9]

The four categories differ in their potential effectiveness and cost. To modify a legacy item would probably be cheaper than developing a completely new one, but a new item would most likely be more effective and more attune to current requirements than the modified legacy item. Roughly and generally speaking, the main thrust of the AoA process is to analyze the trade-offs between these two contrasting aspects—effectiveness and cost.

2.2 Criteria

The merits and weaknesses of the alternatives are evaluated by criteria, which represent various aspects related to the operation, functionality, and reliability of the alternatives, and the cost factors associated with acquiring, handling, and maintaining them. In general, the set of criteria for evaluating defense (physical) acquisition items such as weapons, sensors, and platforms is divided into four subsets:

- Effectiveness (operational performance)
- Operationability, Reliability, Maintainability, and Logistics (ORML)
- Cost and Schedule
- Risk



While performance and effectiveness are measured by specially constructed measures of effectiveness (MOE), and cost is typically measured in money spent (and/or to be spent), ORML criteria are measured by both MOEs and cost factors.

Performance and Effectiveness

The main driver for selecting an alternative for acquisition is its effectiveness – its usefulness in satisfying performance requirements. The terms *effectiveness* and *performance* may be elusive and may mean different things. Measuring effectiveness of an alternative is probably the most challenging part of an AoA study. To demonstrate the complexity of this challenge, consider the following simple (in fact, simplistic) example.

The requirement is for an anti-air (AA) weapon, and the only two criteria are fire-rate and single-shot kill probability (SSKP). There are two alternative weapons for consideration. Weapon A has a higher fire rate than Weapon B, but smaller SSKP than Weapon B. Which weapon is more effective? Weapon A can deliver higher “quantity” of shots while Weapon B has a better “quality” per shot.

One way of measuring the (relative) effectiveness of the two AA weapons is to determine a tactical or operational objective (e.g., maximize number of target killed), determine an appropriate MOE (e.g., expected number of killed targets within a certain time period), and construct a model (analytic or simulation) that calculates the values of the MOE for the two weapons. Another way to determine the relative effectiveness of the two weapons is to treat each attribute—fire-rate and SSKP—as separate criteria, give a score to each weapon with respect to each criterion, and then combine the scores of the two criteria, via a weighted combination, into a single score for each weapon.

The first approach could be considered “objective” in the sense that there is a quantitative model that bridges between the data and assumptions and the final evaluation of the two weapons. The second approach is “subjective” in the sense that stakeholders and/or subject-matter experts must provide their personal input in determining the scores of the alternatives with respect to the criteria and the weights of the criteria.

In reality, and unlike the above example, *effectiveness* has more than two aspects, and thus measuring effectiveness becomes much more challenging. Ideally, there would



be one measurable objective (e.g., maximize probability of winning the engagement, battle, or campaign) that encompasses all relevant operational aspects of the item and the scenario in which it is to be implemented (see Section 2.3). The measurable objective would be formalized as an MOE, which could be reliably computed in a comprehensive model. Unfortunately, this ideal setting seldom occurs; either there are multiple objectives or the scenario and the role of the item in it are so complex that no model can reliably capture all the salient aspects.

The bottom line is that in reality, effectiveness in an AoA is evaluated by a mixture of the two approaches—the analytic “objective” approach and the opinion, or experience-based, “subjective” approach. The goal is to enhance, as much as possible, the analytic side and thus minimize the possible biases and disagreements (see Section 2.4) that may be generated from the subjective approach.

Operationability, Reliability, Maintainability, and Logistics (ORML)

During its course of operations an item must be operated (or controlled) by qualified persons, professionally maintained, and regularly serviced and resupplied. These requirements result in operational, as well as economic, implications. Obviously, *ceteris paribus*, an item that is more reliable and requires less maintenance, less-qualified operators, and lighter logistic burden is preferred to an item that is rated worse on any of these aspects. The question is that of tradeoffs—how much effectiveness one would be willing to sacrifice for a simpler, more robust, and lower-maintenance system?

Operationability is a criterion that reflects two salient aspects of a newly acquired item: (a) compatibility with existing systems currently in use, with which the new item has to interact, and (b) human-system integration (HSI). New radar must interact with existing sensors, command, and control systems and weapons, and therefore must be compatible with them. However, measuring compatibility is challenging; there is no natural MOE that could be defined and objectively evaluated for measuring how well a certain alternative item interacts with current systems. This is a subjective criterion that must be evaluated qualitatively by subject-matter experts (SMEs). Similar restrictions also apply to HSI, the other part of operationability. While, in principle, one could determine the number of operators required for operating the item, measure (say, by length of service, grade, etc.)



the qualifications needed, and estimate the length of the training period, in practice it would be very difficult to define an MOE (or a set thereof) that can adequately represent this criterion. Here, once again, evaluating this criterion is done by qualitative input obtained from SMEs.

Reliability affects the readiness of the item. The more reliable an item is, the less frequent are the times when it is unexpectedly down. This criterion is quantitative and is typically measured by the mean time between failures (MTBF). While measuring the MTBF of an existing system is a relatively straightforward statistical task, estimating the value of this criterion for items in a design or development phases is challenging because of lack of statistical data. Thus, reliability estimates must depend on engineering-based projections based on the item's design and the technical specifications of its components. In many cases, these estimates are provided by the vendors of the items, in which case the projected reliability values must be taken cautiously.

Maintainability is an attribute that describes the technical and physical burden associated with an item. Arguably, a modular item that requires a "plug-and-play" type service is more maintainable than an item that comprises hard-wired components, which, when failed, result in the need for system-wide service. Similar to reliability, maintainability of an item could be measured by MOEs such as average mean service time over all components of the system, or the maximum mean service time, or other statistical measures of repair and maintenance services. The same challenges that apply to the reliability criterion, when an item is still in the design or development phases, apply here too. Maintainability could be considered a fully-quantitative criterion only for existing items, which have accumulated enough maintenance experience and data. Otherwise, maintainability is evaluated by SMEs.

Logistics refers to the operations and cost aspects related to the transshipment of items, the supply chain of consumables (e.g., ammunition, fuel), and repair parts needed for their operation and maintenance. There are typically two logistic aspects associated with an item: (a) the physical infrastructure needed for storage and maintenance of the item and (b) transportation and handling equipment for transporting the item and its required supplies. For example, transporting fuel requires specially designated tankers.



Certain items may also impose logistic constraints. For example, Vertical Launching Systems (VLS) missiles currently cannot be resupplied at sea. In order to replenish this type of ammunition warships must return to port.

Cost and Schedule

While end-users of a military item (commanders, combat developers, operations officers, etc.) are mostly concerned with effectiveness and operationability of the item, program managers and budget officials may be mostly concerned with its overall cost (see Stakeholders section below). Cost comprises several expenditures that vary in their nature (e.g., R&D, production, life-cycle), the time horizon during which they are to be realized, and the certainty regarding their monetary size. Arguably, costs related to future expenditures (e.g., maintenance) are more uncertain than the R&D cost for an item in an advanced development stage, or than the purchasing price for an off-the-shelf item. The cost criterion can be broken down to sub-criteria representing its various components in order to reflect preferences of immediate versus future expenditures.

The *Schedule* criterion reflects the time-scale maturity of an alternative and the projected time to deployment. The urgency of the item affects the weight of this criterion. If schedule and resources are a concern, the guidance can define one set of alternatives that are analyzed, and another group (or further stratification of the initial group) that should be analyzed as resources allow. This criterion is for assessing the likelihood of completing development, integration, and operational testing activities on schedule and within budget. Where significant risks are identified, the assessment may outline practical mitigation strategies to minimize impact on the operational capability provided to the warfighter, and if applicable, possible workarounds in the event the risks are realized.

Risk

The most complex and elusive criterion is risk, and, arguably, it mostly applies to items that are still at the early or middle stages of development. The risk could include delays in development, production, and assimilation schedules and the cost running over budget (See *Cost and Schedule* above). Alternatives, which are already existing items, or very close to maturity, have relatively little or no risk regarding their availability at the



time when they are needed. The only uncertainty is usually associated with production and funding. For less mature alternatives, the more technologically challenged the item, the higher the probability that something will go wrong. The problem is that it is extremely difficult to estimate this probability. Therefore, this criterion is essentially subjective, where risk assessments are mostly based on inputs from subject matter experts or qualitative projections based on data from past similar experiences. Under this criterion we also include the sensitivity of each alternative to possible changes in key assumptions or variables.

2.3 Scenarios

Scenarios may be considered as Uber criteria. They form the settings in which the alternatives are evaluated with respect to the “regular” criteria described in Section 2.2. There are two types of scenarios to be considered in an AoA study. The first type refers to the operational setting in which the item is designated to operate. Military conflict scenarios—and in particular, combat scenarios—are used for in-context evaluation of the effectiveness, operationability, and logistics of items such as weapons, C2 systems, sensors, and other defense- and military-related items. For example, the level, or weight, of importance of the range of a sensor—an Effectiveness criterion—may depend on the typical detection ranges applicable in a certain operational scenario. The importance of the robustness of a vehicle to road conditions may depend on the typical terrain in a scenario. Thus, the designated operational setting of an item is important for evaluating the item’s potential effectiveness. An alternative that performs well over a wide range of plausible scenarios may be preferred to an alternative that performs very well on limited operational settings but poorly on other feasible settings.

The second type of scenarios applies to AoA of items that do not yet exist and are in various stages of research and/or development phases. These scenarios describe economic, social, political, and technological factors that may affect the risk associated with selecting a certain alternative. For example, if a certain alternative requires a considerable R&D effort, the risk, and perhaps the cost criteria associated with that alternative will be impacted by the availability of economic resources and technological capabilities.



Both types of scenarios incorporate a fair amount of uncertainty that must be factored in the AoA study. The way scenarios are incorporated in an AoA study is discussed in more detail in the modeling part described in Section 4.

2.4 Stakeholders

As much as it is a technical and analytical process, acquisition AoA is also a social phenomenon involving a plethora of stakeholders who may represent different interests, viewpoints, agendas, and goals. For example, combatants (the future users of the item) may focus on the effectiveness of the item and its compatibility with existing combat system currently in use. Combat developers may look at a much wider picture and will be concerned with issues of force structure and other strategic considerations. Technical experts will focus on the scientific and engineering aspects, and, in particular, on potential technological challenges that may affect the risk criterion. Finally, budget officials will naturally focus on the programmatic aspects associated with the developing, production, operation, and maintenance of the item. In other words, the cost criterion plays a major part for these stakeholders.

Because the AoA process is complex and multidimensional, and some criteria (dimensions) may be conflicting, it is important to select a balanced mix of stakeholders for the study—representing all the aforementioned groups of decision makers and experts who represent different aspects of the decision problem.



3. The Set of Criteria

The criteria are the touchstones that determine the in-context value of an alternative. Obviously, the goal is to select the alternative with the highest overall value when all relevant criteria are considered. The set of criteria should adhere to some structural, as well as content, rules and properties, which are described in the following sections.

3.1 Criteria Tree

It is convenient to view the set of criteria as a hierarchical structure. This view is not new; it is manifested, for example, in the Analytic Hierarchy Process (AHP) [10], which is used by the DoD. The idea is to break down the main four criteria—(1) effectiveness; (2) operationability, reliability, maintainability, and logistics (ORML); (3) cost; and (4) risk—into sub-criteria, sub-sub-criteria, and so on. This breakdown induces a tree structure whose leaves (lowest hierarchy) are criteria that can either be measured by Measures of Performance (MOPs) or Measures of Effectiveness (MOEs), or can meaningfully be evaluated qualitatively by subject matter experts (SMEs). The aforementioned four criteria constitute the first layer of the criteria tree, as shown in Figure 1.

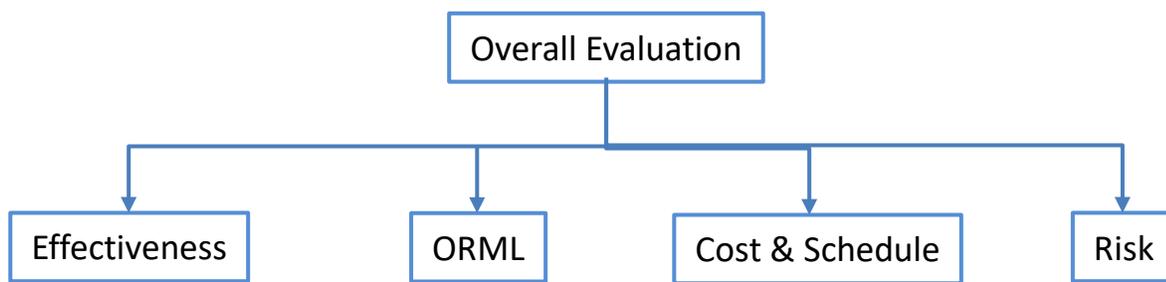


Figure 1. First Layer of a Criteria Tree

If the item to be selected is, for example, some kind of a ground fighting vehicle (e.g., a tank), then possible second and third layers of sub-criteria, which evolve from the effectiveness criterion, are shown in Figure 2.

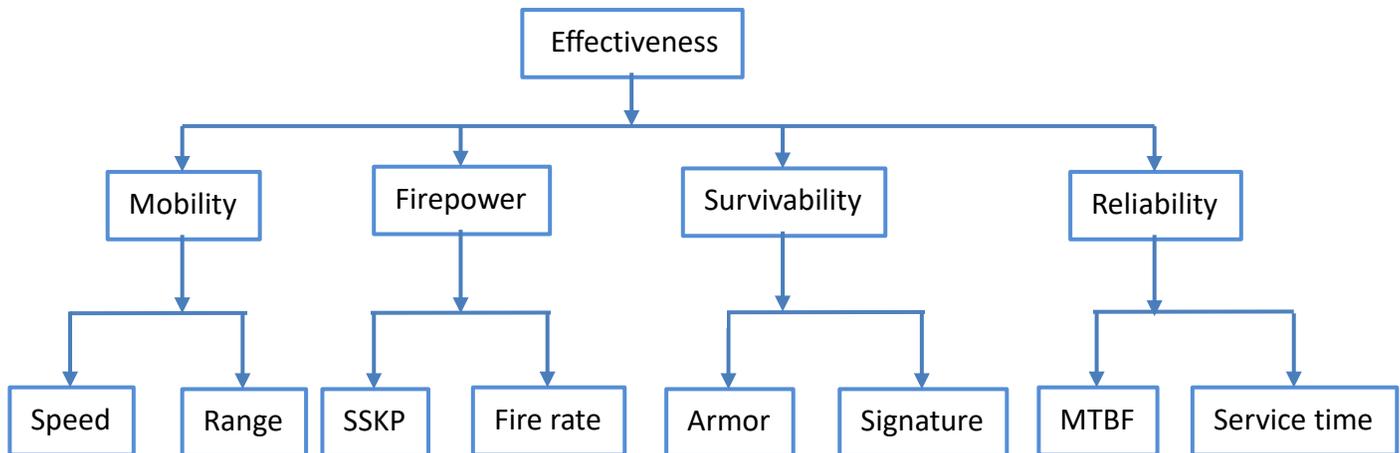


Figure 2. Second and Third Layers of a Criteria Tree

The third layer can be further broken down. For example, *speed* could be divided according to the type of terrain, and *SSKP* may be separated into *day-SSKP* and *night-SSKP*. And so on.

3.2 Criteria Properties

A proper design of the criteria tree is crucial for the success of an AoA project. Specifically, the number of layers in the criteria tree and the granularity of each layer depend on the context and thrust of the analysis and on the complexity of the parent criterion in a higher layer. On the one hand, it is important to include all relevant sub-criteria that affect the parent criterion. On the other hand, we need to avoid over-cluttering the criteria tree such that it remains as manageable as possible. Keeney and Raiffa, in their seminal work [11], suggested some rules or properties that should guide the way criteria are selected for the analysis. In particular, the set of criteria must be *complete* in the sense that the “leaves” of the criteria tree (the end criteria at the lowest layer) cover all the aspects affecting the choice of the item. The criteria must also be *operational*—it must be relevant to the decision problem and meaningful to the decision makers. Another important rule is to avoid *redundancy* that can lead to the undesirable effect of double

counting. For example, the criterion “range of an aerial platform” may be redundant in the presence of the criterion “endurance of an aerial platform.” Finally, as mentioned above, the set of (end) criteria must be as small as possible, notwithstanding the other properties.

It is noted that in some cases breaking down a criterion to more refined sub-criteria (in a lower layer of the criteria tree) may be counterproductive when the sub-criteria are interdependent. Two criteria are dependent if the importance or *weight* (see Section 3.3 below) of one criterion is affected by the evaluation of the alternative with respect to the other criterion. For example, the speed and maneuverability of a fighter aircraft might be dependent; if the aircraft is slow, the maneuverability may be more important than if the aircraft is fast. In that case, the two sub-criteria may be combined into a single criterion such as *flight capability*.

3.3 Weights of Criteria

Different criteria may have different levels of importance, or different *weights*. An important fact to remember is that these weights are subjective. There is no scientific method that could measure the “true” weight of a criterion. Different stakeholders may have different opinions regarding the impact a certain criterion has on the overall value of an alternative. Moreover, the weight of a criterion may also depend on the scenario; certain capability of an item may vary according to the scenario in which the item is to be employed. For example, the importance of the criterion “EO Signature” of a platform depends on the detection capabilities of the adversary in a conflict scenario. Absent such capabilities, the weight of this criterion will most likely be quite low. Another example is the reliability of equipment. If the system has large redundancy with respect to the availability of this equipment, then the weight of the *reliability* criterion may be lower than in the case were the system relies on a single availability of that equipment. Also, the economic, political, social, and technological scenario may affect the weight to be assigned to the risk-related criteria. In Section 4 we describe a method for eliciting weights that take into account the aforementioned factors: multiple stakeholders and multiple scenarios.



THIS PAGE LEFT INTENTIONALLY BLANK



4. Methodology

The AoA process is about comparing the *values* of the alternatives. The best alternative—the one to be selected for acquisition—is the alternative that provides the highest overall value. But, how can one combine multiple criteria and opinions into a single value? What is the scalar function that translates measurements and evaluations of the alternatives with respect to the various criteria and evaluations of criteria weights into a single value that can be compared among alternative items? The problem is exacerbated in the presence of multiple stakeholders who may provide a plethora of opinions and multiple scenarios that may result in different in-context evaluations.

We start with the basic construct, which is common in many decision settings—a linear value function (see e.g., [10]). Simply stated, if w_j denotes the weight of criterion j and v_{ij} is the value of alternative i with respect to criterion j , $i = 1, \dots, I, j = 1, \dots, J$, then the overall value of alternative i , V_i^* , which is to be compared with the overall values of the other alternatives, is given by

$$V_i^* = \sum_{j=1}^J w_j v_{ij}, \quad i = 1, \dots, I \quad (1).$$

The alternative with the highest V_i^* value is the most likely alternative to be selected. Note that we state “Most likely” and not “definitely.” This distinction is discussed later on in Sections 5 and 6.

As mentioned earlier, the challenges are to determine, in a most acceptable way, the values of w_j , $j = 1, \dots, J$, and v_{ij} , $i = 1, \dots, I, j = 1, \dots, J$, taking into consideration the presence of multiple stakeholders and multiple scenarios. We will construct our value function step-by-step, starting with determining the values of the weights w_j .

4.1 Determining Criteria Weights for a Certain Scenario

Consider a certain reference scenario s . Assuming this scenario is realized, we wish to elicit from R stakeholders criteria weights w_{js} , $j = 1, \dots, J$, that (a) reflect the relative



importance of the various criteria if scenario s prevails, and (b) represent a consensus among the stakeholders regarding these weights. An efficient and effective way to elicit preferences from decision-makers is through ratio-scale matrices, similarly to the setup used in AHP (see [10]). The idea is as follows: each stakeholder $r, r = 1, \dots, R$, is asked to compare two criteria weights, say w_{js} and w_{ks} , with respect to scenario s . In other words, the stakeholder provides an assessment regarding the extent one criteria is more (or less) important than the other. The comparison is in terms of the ratio between the two weights.

That is, p_{jks}^r is the assessment of stakeholder r regarding the ratio $\frac{w_{js}}{w_{ks}}$. Different stakeholders may have different opinions regarding the very same ratio. In other words, for two stakeholders r and r' , we may have $p_{jks}^r \neq p_{jks}^{r'}$. A natural way to mathematically resolve such discrepancies is using least squares. The same way least squares are used to fit a “consensus” line among a clutter of points in statistical linear regression, we can derive a consensus set of weights by minimizing the L_2 distance. Formally, we solve the following non-linear optimization problem:

$$\begin{aligned} & \text{Min} \sum_{j < k} \sum_{r=1}^R \left(\frac{w_{js}}{w_{ks}} - p_{jks}^r \right)^2 & (2) \\ & \text{st} \\ & \sum_{j=1}^J w_{js} = 1, \quad w_{js} \geq 0. \end{aligned}$$

The objective function is separable, quadratic, and therefore the optimization problem is easily solvable for real-size problems by tools as simple as the MS Excel Solver. The constraint is just a normalization of the criteria weights, which facilitates simpler computations down the road.

The optimal solution of Problem (2) is a vector $W_s^* = (w_{1s}^*, \dots, w_{js}^*)$ of criteria weights that represent an L_2 consensus regarding the criteria importance in the presence of scenario s .



We solve Problem (2) S times—once for each possible scenario. For brevity and simplification, we drop the $*$ sign from future notation.

The model in (2) described above for criteria weights can obviously be applied, sequentially, to the different levels of the criteria tree (see Section 3.1 above). For each master criterion at level l , we solve problem (2) for the “child” criteria at level $l + 1$. The weight of the end criterion at the bottom level is the product of the criteria weights leading to that criterion. For example, considering Figures 1 and 2 in Section 3.1, we first solve problem (2) for *effectiveness*, *ORML*, *cost*, and *risk*. Next, for the master criterion *effectiveness*, we solve (2) for *mobility*, *firepower*, *survivability*, and *reliability*. Similarly, Problem (2) is solved for the childes (if any) of *ORML*, *cost*, and *risk*. Finally, we solve (2) for the lowest level (e.g., *speed* and *range* for *mobility*, *SSKP* and *fire-rate* for *firepower*, etc.). The weight of the end-criterion in the value function (1), say *speed*, is the product

$$W_{Effectiveness,s} \times W_{Mobility,s} \times W_{Speed,s}.$$

4.2 Determining Alternatives' Values for a Certain Criterion and Scenario

Once again, we consider a certain scenario s . Let us also consider a certain criterion j . Similarly to the way criteria's relative weights are elicited from stakeholders, the objective here is to obtain the ratio figure $d_{ij_s}^r$ that represents stakeholder's r opinion regarding the ratio between the value of alternative i and alternative l with respect to criterion j , in the presence of scenario s . Similarly to (2), we solve now

$$\begin{aligned} & \text{Min} \sum_{i < l} \sum_{r=1}^R \left(\frac{v_{ijs}}{v_{ljs}} - d_{ij_s}^r \right)^2 & (3) \\ & \text{st} \\ & \sum_{i=1}^I v_{ijs} = 1, \quad v_{ijs} \geq 0. \end{aligned}$$

Problem (3) has an identical structure as Problem (2). Here, the optimal values $V_{j_s}^* = (v_{1j_s}^*, \dots, v_{Ij_s}^*)$ are the mathematical consensus values of the alternatives with respect to criterion j , under the assumption of scenario s . Problem (3) is solved $J \times S$



times, once for each criterion and each scenario. As before, for brevity and simplification, we drop the * sign from future notation.

Note that in both Problems (2) and (3) we assume a homogeneous or “democratic” set of stakeholders; no stakeholder’s opinion is considered more influential, or with higher weight, than others. If this is not the case, and certain stakeholders’ opinions weigh more than others, then the objective functions in (2) and (3) are weighted accordingly with stakeholders’ r -indexed weights. The problems are still easy to solve.

4.3 The Alternative’s Value Function in a Scenario

Following the solutions of problems (2) and (3) we can compute the value of an alternative, with respect to a certain scenario. This value reflects the consensus weights of the criteria $W_s = (w_{1s}, \dots, w_{js})$, and the consensus (relative) values of the alternative $V_{js} = (v_{1js}, \dots, v_{ijs})$. Formally, the consensus overall value of alternative i in scenario s is given by

$$\bar{v}_{is} = \sum_{j=1}^J w_{js} v_{ijs} \quad (4)$$

In other words, if an “Oracle” could tell us the scenario s to be unfolded, then the alternative i to be considered for selection is the one for which \bar{v}_{is} is maximized. Absent such Oracle, the probabilities of the various scenarios must be taken into consideration when trying to identify the best alternative.

4.4 Scenarios’ Probabilities

Nobody knows for sure which of the possible scenarios will actually be realized. Different stakeholders may have different opinions about the likelihood of the various scenarios. The combined assessment of scenarios’ probabilities is obtained using the same methodology as in Problems (2) and (3).

Let a_{st}^r denote the assessment of stakeholder r about the relative likelihood of scenarios s and t . That is, a_{st}^r is the subjective opinion of stakeholder r regarding the extent scenario s is more (or less) likely than scenario t . The consensus probabilities



$q_s, s=1, \dots, S$, of the various scenarios is obtained as the solution of the quadratic optimization problem,

$$\begin{aligned} & \text{Min} \sum_{s < t} \sum_{r=1}^R \left(\frac{q_s}{q_t} - a_{st}^r \right)^2 & (5) \\ & \text{st} \\ & \sum_{s=1}^S q_s = 1, \quad q_s \geq 0. \end{aligned}$$

The optimal solution $Q^* = (q_1^*, \dots, q_S^*)$ is the consensus probability distribution of the scenarios. As before, we drop the * sign from future notation.



THIS PAGE LEFT INTENTIONALLY BLANK



5. Selecting the Winning Alternative

Following the operations described in Section 4, the AoA team has an initial set of parameters that reflect the stakeholders' L₂-consensus regarding (a) the weight of criteria in each scenario, (b) the relative value of each alternative with respect to each criterion in each scenario, and (c) the (subjective) probabilities of the scenarios.

Recall from Section 4.3 that for each scenario s we have now a calculated value \bar{v}_{is} for each alternative i . This value represents the L₂-consensus outcome of the stakeholders' group decision process with respect to the relative standing of alternative i , if scenario s is realized. The L₂-consensus about the likelihood of scenario s is q_s . Thus, we have now a (subjective) probability distribution of alternatives values over scenarios where each value \bar{v}_{is} is associated with a probability q_s .

There are several ways to proceed from this point and identify the alternative that is most likely to be the best among the set of I alternatives. The most natural measure is the expected value where the “winning” alternative is alternative i for which $\sum_{s=1}^S q_s \bar{v}_{is}$ is maximized. Here we choose the alternative that “on-average” over the possible scenarios produces the highest relative value. This linear measure is quite common and easy to explain to decision-makers, but it is not always the right yardstick for choosing an alternative, in particular when the specific likelihoods of scenarios are to be looked at in more detail.

Another possible measure is the mode of the distribution; we simply choose the alternative that performs the best with respect to the *most likely* scenario. That is, if $s' = \arg \max q_s$ then the selected alternative i is the one for which $\bar{v}_{is'}$ is maximal. This measure is appropriate if there is one scenario that stands out as very likely—much more than any other scenario. If the induced (subjective) entropy of the scenarios, as implied from the stakeholders' projections, is high, then obviously the mode measure will be inappropriate.



Lastly, and probably most appropriately, it would be better to select an alternative that is *good* over a large set of scenarios than an alternative that is *excellent* over a smaller set of scenarios. The goal here is to seek robustness in the choice of the winning alternative. The idea is as follows.

First, we set a probability threshold, which represents the level of confidence, with respect to the realized scenario, which we wish to associate with the winning alternative. Suppose this probability level is α . Reasonable values of α are in the range 0.6–0.9. Next, we generate all the minimal subsets of scenarios whose combined probabilities are at least α . For each such subset, we identify the alternative(s) for which the **minimum** value across the scenarios in the subset is **maximal**. Formally, let T_1, \dots, T_M denote the set of all the α -valued subsets of scenarios. Each subset T_m comprises scenarios with combined probabilities of at least α , and any scenario removed from that set reduces the combined probabilities to less than α (hence, minimal subsets).

Suppose a certain subset $T_m = \{s_1, \dots, s_n\}$. Note that we identify a scenario by its index, and that $\sum_{k=1}^n q_{s_k} \geq \alpha$. Define $v_i(T_m) = \text{Min}_{s \in T_m} v_s$ and $i_m = \arg \max \{v_1(T_m), \dots, v_I(T_m)\}$. Alternative i_m is the max-min alternative for scenarios subset T_m . In other words, alternative i_m provides the highest *guaranteed* value among all alternatives if one of the scenarios in scenario set T_m is realized. Finally, $i^* = \arg \max \{v_{i_1}, \dots, v_{i_n}\}$ is the alternative that has the highest value with probability of at least α . Obviously, v_{i^*} is monotone non-increasing in α ; the higher the required probability threshold the smaller the assured alternative value. To demonstrate this procedure, consider the following example:

- Three scenarios, A, B, and C, with probabilities 0.3, 0.3, and 0.4, respectively
- There are 3 alternatives
- The overall values \bar{v}_{is} of the (alternative x scenario) combinations are shown in the Table 1.



Table 1. Values of Alternatives

Scenario Alternative	A (0.3)	B (0.3)	C (0.4)
1	0.7	0.5	0.9
2	0.6	0.8	0.6
3	0.9	0.4	0.5

Suppose $\alpha = 0.6$, which implies that we have here three subsets of scenarios that satisfy the minimum probability threshold requirement: $T_1 = \{A, B\}$ with probability $0.3 + 0.3 = 0.6$, $T_2 = \{A, C\}$ with probability $0.3 + 0.4 = 0.7$, and $T_3 = \{B, C\}$ with probability $0.3 + 0.4 = 0.7$. For the first scenario set we have: $v_1(T_1) = 0.5$, $v_2(T_1) = 0.6$, $v_3(T_1) = 0.4$, and therefore the max-min alternative i_1 is alternative 2 with value 0.6. For the second scenario set we have: $v_1(T_2) = 0.7$, $v_2(T_2) = 0.6$, $v_3(T_2) = 0.5$, and therefore the max-min alternative i_2 is alternative 1 with value 0.7. For the third scenario set we have: $v_1(T_3) = 0.5$, $v_2(T_3) = 0.6$, $v_3(T_3) = 0.4$, and therefore the max-min alternative i_3 is alternative 2 with value 0.6. Thus, alternatives 1 and 2 are candidates for selection. But the maximum value over the eligible scenario sets is 0.7 and is obtained by alternative 1. Therefore, at confidence level of at least 0.6 the highest-valued alternative is alternative 1.

Notice how the likelihoods of the scenarios affect the choice of alternatives. If, instead of the probabilities values in Table 1, the scenario probabilities were 0.6, 0.2, and 0.2 for scenarios A, B, and C, respectively, then it is easily seen that alternative 3 becomes the most preferred one with value 0.9. Going back to the original probabilities, if the threshold α is now 0.8, then we only have one subset (the complete set of scenarios) and the max-min alternative is alternative 2 with min value of 0.6.

To summarize, this quantile-type approach is both flexible, in the sense that one could choose the confidence level for selecting the best alternative, and robust by adopting the max-min measure of alternatives' values. This approach selects an alternative that is good over a wide range of possible scenarios instead of an alternative that is excellent in only limited number of situations.



THIS PAGE LEFT INTENTIONALLY BLANK



6. Implementation

In the last two sections, we described a formal process for conducting AoA by a group of stakeholders who may have different perspectives and opinions regarding the subject matter. Disagreements and inconsistencies in preferences and assessments may occur with respect to criteria weights, alternative valuations, and scenario likelihoods. The proposed group-decision model produces a consensus rating of the alternatives based on minimizing disagreements in the L_2 metric sense.

While the model is transparent and relatively simple to implement in a spreadsheet, it should not be considered as a “black box” that automatically produces a winning alternative based on stakeholders’ and analysts’ inputs. The “winning” alternative that emerges from the model may not be the final choice in situations where the value(s) of the runner-up(s) is (are) not significantly different from the winner’s value. The model is a technical tool that, following a properly designed sensitivity analysis, can help guide the AoA process towards a robust decision. A possible paradigm for conducting a well-structured analysis of acquisition alternatives is as follows:

Step 1: Establish the AoA team that must comprise a balanced group of stakeholders (e.g., field commanders and end-users), decision makers (e.g., budget managers, defense officials), and analysts (e.g., engineers, cost-estimators, operations-research analysts, and other subject-matter experts).

Step 2: The AoA team reviews documents describing operational setting, requirements, and capability gaps. An initial set of possible acquisition alternatives is generated. The analysts in the team start gathering more detailed data and information about the operational setting and the possible alternatives.

Step 3: Non-starter alternatives are identified and removed from consideration. Such alternatives are items that are rejected up front because of reasons such as: not meeting minimum capability thresholds, being too costly, or being based on immature technologies. The team defines the sets of alternatives, criteria, and scenarios. This step comprises open discussions that set the stage for the detailed analysis to follow.



Step 4: Each member r in the AoA team provides her/his estimates for: p_{jks}^r (see Section 4.1), d_{iljs}^r (see Section 4.2) and a_{st}^r (see Section 4.4). This step includes also operations-research and cost-estimation analyses, which provide valuable input to the AoA team.

Step 5: Model implementation on data gathered in Step 4. Output: set of alternative ratings.

Step 6: Discussion on the model results (alternative ratings) and performing sensitivity analysis on all three factors: criteria weights, alternatives' values, and scenarios' probabilities. Step 5 may be repeated several times based on the discussions in this step.

We can see that the model described in Section 4 acts as a decision aid and facilitator for discussions among the team members rather than an "Oracle" that crunches numbers and provides a "solution."



References

- [1] Department of Defense (DoD). (2017, August 10). *Operation of the defense acquisition system* (DoD Instruction 5000.02). Washington, DC: Author.
- [2] RAND Project Air Force. (2006). *Analysis of alternatives (AoA) for KC-135 recapitalization*. Executive summary. Santa Monica, CA: RAND Corporation.
- [3] TRADOC Analysis Center. (2011). *Joint light tactical vehicle analysis of alternatives final report* (TRAC-F-TR-11-025).
- [4] Department of the Navy. (2006). *Seabase to shore connector analysis of alternatives final report*. Reston, VA: WBB Consulting.
- [5] Souders, R. M., Schulze, S., Ginburg, Y., & Goetke, J. (2004). *MPF(F) analysis of alternatives: Final summary report* (CNA, CNR D0009814.A2/Final).
- [6] Smith, M. W., Ritz, C., Smith, M., & Braun, C. (2011). *T-AO(X) analysis of alternatives final report* (CNA, CRM D0026011.A2/Final).
- [7] RAND. (2016). *Maritime tactical command and control analysis of alternatives*. Santa Monica, CA: Author.
- [8] MITRE. (n.d.). Analyses of alternatives. In *Systems engineering guide*. Retrieved on April 23, 2018, from <https://www.mitre.org/publications/systems-engineering-guide/acquisition-systems-engineering/acquisition-program-planning/performing-analyses-of-alternatives>
- [9] U.S. Air Force, Office of Aerospace Studies. (2016, July 6). *Analysis of alternatives handbook*. Retrieved from <http://afacpo.com/AQDocs/AoAHandbook.pdf>
- [10] Saaty, T. L. (1980). *The analytic hierarchy process*. New York, NY: McGraw Hill.
- [11] Keeney, R. L., & Raiffa H. (1976). *Decisions with multiple objectives*. Hoboken, NJ: J. Wiley & Sons.





ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CA 93943

www.acquisitionresearch.net