2-27-2013

A Cost-Based Decision Tool for Valuing DoD System Design Options

Erin T. Ryan
Air Force Institute of Technology

David R. Jacques
Air Force Institute of Technology

Jonathan D. Ritschel
Air Force Institute of Technology

John M. Colombi
Air Force Institute of Technology

Available online at www.sciencedirect.com

SciVerse ScienceDirect

Procedia Computer Science 16 (2013) 1180 – 1189

Procedia
Computer Science

Conference on Systems Engineering Research (CSER'13)
Eds.: C.J.J. Paredis, C. Bishop, D. Bodner, Georgia Institute of Technology, Atlanta, GA, March 19-22, 2013.

A Cost-Based Decision Tool for Valuing DoD System Design Options

E. Ryan, D. Jacques, J. Ritschel, J. Colombi

Department of Systems Engineering and Management
Air Force Institute of Technology
2950 Hobson Way, WPAFB, OH 45433
erin.ryan@afit.edu

Abstract

The DoD has frequently demonstrated its ability to procure phenomenal systems; however, these accomplishments are often tarnished by substantial cost and schedule overruns. While defense acquisition policies are continually being revised to address these perennial problems, many believe that a more fundamental source of these overruns is the lack of flexibility in the systems being developed, which tend to preclude effective responses to unexpected events. However, providing justification to invest in flexibility is a tough sell when the measure of value is a military capability or political outcome, as there is no extant method to demonstrate the potential return on investment. This paper introduces a decision tool for valuing the inherent ability of different systems or designs to respond to uncertainty. The proposed tool is essentially a modification of the current life cycle cost model and is premised on the notion that the need for capability changes in a system arises in a stochastic manner that can be incorporated into a continually updated, expected value model presented in terms of total life cycle cost. The cost-based decision tool presented here quantifies the ability of competing designs to respond to these capability changes via a cumulative distribution function (CDF). The design with the most favorable CDF (i.e., the one that is most likely to meet the most likely set of requirements at the lowest expected value curve of life cycle cost) is deemed to be the "best" design.

Selection and/or peer-review under responsibility of Georgia Institute of Technology

Keywords: Life Cycle Cost, Decision Analysis, Stochastic Model, Uncertainty, Affordability

1. Introduction and Motivation

Perennial cost and schedule overruns have become the norm for DoD programs. To those familiar with the history of defense acquisition, these systemic deficiencies are widely known, as is the standard response from the Pentagon. If the past is any indication of the future, then we will soon see another acquisition reform effort spawned and promulgated with the expressed intent of reducing monetary waste, accelerating acquisition timelines, or improving overall mission responsiveness. Sadly, the reform efforts are not likely to make a difference. This observation is not meant to disparage the various well-intentioned reform efforts and the dedicated professionals that create and implement them; the point is, rather, that the desired improvements are seldom, if ever, realized [1,2,3].

Increasingly, DoD leadership is less likely to attribute these cost and schedule overruns to flaws in the acquisition process, but instead to the lack of flexibility in the systems being developed. If systems can be designed in such a way that they are able to more readily respond to various sources of change, then it stands to reason that when uncertainties become realities, the impact to the program will be lessened. Among DoD policy makers and acquisition professionals, there is broad consensus that greater flexibility in weapon systems is a desirable goal. And yet, it is rarely achieved.

The problem is that flexibility necessarily incurs additional investment costs. The most obvious is the direct cost

Selection and/or peer-review under responsibility of Georgia Institute of Technology
doi:10.1016/j.procs.2013.01.124

associated with implementing the flexible design, generally related to pre-provisioning the system with nascent capabilities that can be matured to full implementation at a later time. The less obvious cost of flexibility pertains to the tradeoffs that must be made against other performance attributes. The notion of designing to the bleeding edge of performance requirements is antithetical to the aims of flexibility, as it consumes engineering tradespace. An inherently flexible design cannot, axiomatically, achieve the same level of technical performance along every dimension as the performance-optimized design. This "capability cost" of flexibility can serve as an especially strong deterrent in DoD's contemporary, requirement-driven, performance-dominated mindset. In such an environment, the costs associated with more flexible design solutions must be assiduously justified in order to have any hope of being implemented. At present, there is no such method within the DoD.

This paper provides a potential method for valuing "flexible" designs in the form of a rational decision tool for discriminating between competing design options. This tool consists of a top-down, intrinsic value model based on the familiar notion of Life Cycle Cost (LCC). The idea is to refine current LCC calculations to better account for the value of capability opportunities that are likely to arise throughout the life of a program, and that the "best" design that achieves some assured minimum capability is merely the one that is the most likely to be the most cost-effective over its life cycle. The relative measure of cost effectiveness is via a cumulative distribution function of life cycle cost unique to each candidate design.

2. Background

There is a substantial amount of literature on the topic of flexibility, largely focused on how to implement it and measure it. Less common is the question of how to value flexibility. Rarer still is discussion regarding how to value flexibility *within the context of defense systems*. Moreover, the issue of ascribing value to flexible designs is generally regarded as a problem involving *decision making under conditions of uncertainty*. Although there is also ample literature on this subject, the applicability to the DoD is again limited due to a lack of typical value metrics.

2.1. Decision making under uncertainty

Based on the literature, we know that the value of flexibility is positively correlated to uncertainty, such that the greater the uncertainty in the system, the greater the value a flexible design option is likely to have [4,5,6,7]. But if we are to make any headway on quantifying the value of flexibility, we need the ability to make the best decision under conditions of uncertainty.

This type of problem has been studied extensively in economics. One common approach within that community is net present value (NPV) analysis. NPV is a standard method for determining the time value of money. It takes into account the net cash flow at a particular time t, as well as the required rate of return (i.e., *discount rate*). Thus, the expected cash flows are discounted at an interest rate that accounts for the time value of money as well as the project risk. Several studies use NPV as part of their effort to quantify flexibility, including [8,5,9,10].

Another approach is *real options analysis*, which exists at the intersection of value and uncertainty. Economic theory defines a real option as the "right, but not the obligation to take an action at a predetermined cost and at a predetermined time" [6]. Copeland [11] claims that only real options can "provide a theoretically sound tool for valuing" decision flexibility. In a manufacturing application, Ajah [12] states "that the adoption of the real options approach early in the conceptual design process can offer to the designer, extra degrees of freedom of systematically considering and designing system elements."

2.2. The problem with net present value

In general, researchers tend not to be in favor of using NPV for decisions involving flexibility [13]. While NPV is sufficient in cases of "low uncertainty, or [when] you have no scope to change course" [11,14], it is not appropriate for situations involving great uncertainty, as it assumes a predetermined path through an established set of alternatives. This is contrary to the core aim of flexible processes and designs, so a different method is needed—one that can take more (and fewer predetermined) decision options into account [15,16,17].

2.3. The problem with real options

While real options analysis is widely seen as preferable to NPV for uncertainty analysis, it comes with its own set of criticisms, especially if applied to the problem of valuing flexibility in the DoD. For example, the underlying financial model (known as *Black-Scholes*) is generally only valid under certain constraining assumptions, many of which are not likely to be applicable to defense acquisition. One of the model assumptions is that the valued asset must be traded on an "efficient" market, where there is no possibility of arbitrage. While arguably true in the broader capital market, this assumption is not warranted within the DoD monopsony (i.e., single buyer), where markets are often artificial, and far from efficient. Another stipulation of the Black-Scholes model is that the asset must have a price that follows geometric Brownian motion, thus creating a return on the asset that is consistent with a random lognormal distribution. However, random fluctuation of price is a debatable premise in a standard open market, let alone in the cloistered defense industry. Finally, real options in military acquisition programs are likely to be path-dependent and highly interdependent. In both cases, traditional financial options methodologies tend to fail because the underlying stochastic differential equations are not available or simply do not apply [18,17].

2.4. Difficulties associated with valuing military capabilities

In order to make meaningful value judgments, we must establish a utility function that will quantify the value of capability in some ratio-level comparable units. While this is relatively routine for profit-driven commercial systems, it will necessarily be more challenging for military systems, as the utility function will almost certainly not involve a monetizable metric like earnings. Instead, for example, we would need to somehow devise a function (or more likely, a series of functions) for determining the utility of an extremely wide range of military capabilities.

In principle, there is a solution. Under the neoclassic economic definition of value, an item's value can be established by determining a customer's *willingness to pay*. Thus, we can surmise that the value of a particular military capability can be determined by ascertaining the maximum amount the government is willing to give up (of some measureable resource) to obtain the capability (i.e., the value of a given capability to the government = the maximum cost the government is willing to pay for the capability). The devil is in the details, however.

Assigning a numerical value to the right side of this equation (i.e., what the government is willing to pay) is a daunting endeavor. The most obvious approach would be to use the dollar amount budgeted by the government. But this is problematic for a multitude of reasons. Consider that the actual system cost may include a number of other scarce resources (e.g., time, critical skills, facilities) that are not captured in the government budget. Technically, economic cost includes the loss of opportunities as well, so we would also need to account for the cost of losing or vitiating other capabilities by virtue of the fact that we are committing resources to this capability. Once again, though, we would face the dilemma of assigning a value to a capability, with only budgets to guide us, so our original problem is further complicated because it is now recursive.

Finally, even if we were to accept that budgeted costs will be adequate, there is no reason to believe this represents the *maximum* cost the government is willing to pay. Firstly, the government may, in principle, be willing to budget more for a particular capability, but has reason to believe that a lower amount will suffice. The problem is that the government generally establishes its program budgets based on expected actual costs, not the perceived value of the program or resulting capability set. Secondly, budget allocation processes are notoriously volatile, subject to any number of political vagaries that have nothing to do with the merits of a particular program or capability. Thus, one year's total budget allocation for a given program may be substantially different from the next year's allocation for the same program, though there was no change in its perceived value.

3. A More Flexible Approach to Valuing Flexibility

Given the difficulty of establishing the value of military capabilities, it is clear we need a more *flexible* approach to determine the value of flexibility for DoD programs. The question thus arises whether we can establish the merits of a capability without having to explicitly determine its value. This paper argues that this is feasible through a modification to the familiar life cycle cost (LCC) model. The fundamental idea being proposed is to refine current life cycle cost calculations to better account for the value of capability opportunities that are likely to arise throughout the life of a program. Before proceeding to a more comprehensive explanation, however, it may be beneficial to review the salient aspects of DoD's current LCC methodology.

3.1. Life cycle cost

LCC is a systematic accounting approach for aggregating all direct and many in-direct costs for a given system. It includes not just total acquisition costs, but also costs related to operations, maintenance, and disposal. Importantly, LCC also accounts for risks, generally either through sensitivity analyses or through formal quantitative risk analysis [19]. By providing senior decision-makers with their single best source of estimated cost to achieve a given capability, the LCC estimate is often instrumental in determining the ultimate procurement fate of a program.

Formal DoD guidance calls for the LCC to be first accomplished as part of the initial Analysis of Alternatives (AoA) and then updated as part of major milestone decision reviews. Aside from these updates, however, the system LCC is generally a static measurement. When calculated, it provides a "snapshot" estimate of total life cycle cost on the assumption that there will be no deviations from key cost, schedule, and performance parameters, which are collectively referred to as the APB, or Acquisition Program Baseline [19]. Of course, one thing we know with certainty is that there will almost always be deviations from the APB.

While the assumption of a static APB may be unwarranted, programs proceed with it anyway, largely because there must be a foundation upon which to build the cost estimates against, but also because the alternative of trying to account for the non-deterministic uncertainty in precisely *how* the program will deviate from the APB is simply not possible, or at least just too daunting. It can be argued, however, that even though uncertainty is—by definition—not deterministic, it may be possible to employ stochastic probability methods that can yield cost estimates that are likely to be more accurate in the long run [20]. Although establishing the initial models to accomplish this would require significant resource investment, the possibility of more accurate LCC estimates—and the improvement in decision-making that would accompany that—promises an enormous return on such an investment.

3.2. Life cycle cost under uncertainty

Clearly, there is substantial motivation to provide improved LCC estimates, at least to the level required to support decisions considering alternative flexible design options. The notion that this can be done by accounting for random events that affect the system forms the basis of life cycle cost under uncertainty (also referred to as *stochastic life cycle cost*). The idea of applying this strategy to acquiring military systems appears to have been first introduced in two papers related to a DARPA (Defense Advanced Research Projects Agency) satellite program [21,22]. As described by Brown, stochastic life cycle cost is premised on three assertions—

- *The cost to develop, procure, and operate a system with some assured minimum capability over its lifecycle is not a deterministic value.*
- *Instead, this cost can be modeled as a random variable with a probability distribution resulting from a set of uncertainties introduced throughout the system's life.*
- *This random variable metric is a relevant basis for comparison between alternative … design choices.*

Brown is to be commended for introducing this simple, but deceptively powerful, notion of stochastic life cycle cost. However, the initial treatment does not develop the principle fully, nor explore its broader applicability. The type of stochastic events he considers are only those specific events that critically influence the success of a satellite system, i.e., launch failure and on-orbit component failure. Brown explicitly does not consider other aspects of life cycle uncertainty that affect virtually all programs, such as "requirements creep, funding stream volatility, technology development risk, and volatility of demand" [21]. Yet he clearly does recognize that the model could be applied to these other sources of uncertainty, noting that these variables are "left for future analysis." To date, it does not appear that such an analysis has been accomplished by him or others.

Consequently, we propose a research strategy to logically extend this promising technique in a manner that may provide a number of potential benefits over current practices. Specifically, we intend to expand the life cycle cost under uncertainty idea to a robust and comprehensive methodology for effectively valuing system design alternatives.

3.3. Current Expected Value of Life Cycle Cost Curve (CEVLCCC)

To capture the utility of this improved LCC concept, we proffer the term, Current Expected Value Life Cycle Cost Curve, or CEVLCCC (pronounced kev' lik). The name is intended to convey a couple of key distinctions from both the standard LCC and Brown's stochastic LCC. The "Expected Value" phrase discriminates CEVLCCC from the standard LCC as a more probabilistically accurate measurement of system cost; whereas the word "Current" connotes the fact that the CEVLCCC tool is intended to be employed as a continually updated decision analysis tool. The notion that an LCC estimate might be applied dynamically, and at lower levels of system design, is distinct from

Brown's view that the stochastic LCC could only be useful for "preliminary trade space exploration" and not for value determinations "below the architectural level" [22]. Finally, "Curve" denotes that the output is a cumulative distribution function (CDF) of potential costs, not a single point estimate.

Note that under this conception, the "expected value" concept is essentially a penalty that attempts to capture the anticipated cost impacts related to future baseline changes. The more cost-effectively a given design can respond to these changes, the lower the penalty. Given the inherent cost accounting methodology of the CEVLCCC approach, as long as each design is capable of achieving "some assured minimum capability," then the corresponding capabilities and outcomes need not be valued. The relative value can be inferred solely from each design's expected life cycle cost, with the best value presumably the one with the lowest adjusted LCC.

In practice, the proposed methodology is also straightforward, consisting of the following four steps:

1. Establish the System Design Options. First, the user identifies the candidate designs to be evaluated. Each design must be of sufficient maturity that its traditional life cycle cost can be estimated, and cost impacts can be estimated should there be changes related to the assured minimum capability of the system.
2. Construct Time-Phased CDFs. Next, the user creates CDFs to characterize the possibility of change to the assured minimum capability of the system. In practice, this means estimating the probability that the threshold value of existing schedule or technical performance requirements will change at various time points in the future, as well as estimating the probability that specific new requirements will be imposed.
3. Estimate LCC Impacts. The user then estimates LCC impacts associated with the potential changes—as characterized in the PDF—in the assured minimum capability of the system. As part of each estimate, the user specifies a minimum and maximum cost along with an associated confidence ranging from 50 to 90 percent.
4. Select Most Favorable CEVLCCC. The CEVLCCC tool then outputs a probability curve in the form of a CDF of expected life cycle costs associated with each design. If the resulting cost curve of one design is perceived to be more favorable than the other(s), then the user now has a quantitative rationale for choosing among the candidate designs.

The authors have constructed an Excel-based graphical user interface to automate these steps, which we refer to as simply the "CEVLCCC tool."

4. Methodology and Hypothetical Use Case

To appreciate the process and potential utility of the CEVLCCC tool, we illustrate its methodology and application using a hypothetical air superiority stealth fighter program that is considering three competing payload designs. The program wishes to determine which design is likely to be the best value over the program's life cycle. The detailed design differences are not relevant to understanding the principle of the CEVLCCC tool; the reader need only be aware of the basic distinction between each proposed payload design, which is readily inferred from the name of each option. The three options, along with their traditionally estimated life cycle cost, are shown in Table 1. For this exercise, all LCC values are entirely notional, and will be treated as point estimates with no associated error.

Table 1. Payload Designs Options and Estimated Life Cycle Cost

Payload Design Option	Estimated LCC ($)
Small Internal	$1000
Large Internal	$1050
External	$950

Relative to the *Small Internal* payload design, the *Large Internal* design is expected to cost five percent more over its life cycle (mostly due to increased airframe weight) while the *External* design is expected to cost five percent less (due to simpler and proven technology related to externally-mounted armaments). In addition, all three candidate designs are expected to be able to meet or exceed the current threshold values of all current requirements. So under the traditional conception of LCC estimating, and assuming everything else was equal, the payload design above with the lowest estimated LCC (i.e., *External* at $950 million) would typically be the option selected.

If one were to stipulate that the current program baseline will remain fixed (i.e., no changes to the existing set of requirements), it certainly would make sense to choose the *External* design over the other two options. The problem is, of course, that this stipulation is extremely unrealistic. Or, put another way, this approach is too simplistic as *it* *does not account for the value of the flexibility embedded within certain architectural options*. We may not know exactly how or when, but the APB of the system will change, and the fact is that each of the designs has an intrinsically different ability to respond to APB changes. The *External* design may be the best option given a fixed baseline, but what is the best option in the more realistic program future that is characterized by uncertainty? The intent of the CEVLCCC tool is to answer this question in an objective, quantifiable manner.

4.1. Existing (known) requirements

In an actual program, there would likely be a large number of existing schedule and performance requirements that each design would need to be evaluated against as part of a comprehensive CEVLCCC tool analysis. For simplicity, we will consider only the four known requirements shown in Table 2. Notional threshold and objective values are also listed for each of these requirements.

Table 2. Existing Requirements for Stealth Air Superiority Fighter

#	Requirement Description	Threshold	Objective
1	Armament of (X) air-to-air guided missiles	X=4	X=8
2	Nominal front sector radar cross section (RCS) no greater than (X) m^2	X=0.10	X=0.02
3	Top speed of Mach (X)	X=2.0	X=2.5
4	Initial Operating Capability (IOC) within (X) months	X=84	X=60

4.2. New (unknown) requirements

Changes to known requirements are not the only source of uncertainty that should be evaluated. We must also take into account the possibility that new requirements will be levied on the program at some point in the future. There may be several potential new requirements to consider as part of a thorough analysis, but for this simplified scenario, we will evaluate just one potential unknown requirement. As listed in Table 3, the fifth requirement to be evaluated involves the ability of the aircraft to strike ground targets. In other words, although the system does not currently have a formal air-to-ground mission requirement, the program wishes to account for the possibility that this capability will be required at some point in the future.

Table 3. Potential New Requirement for Stealth Air Superiority Fighter

#	Requirement Description	Threshold	Objective
5	Armament of X air-to-ground guided missiles	X=2	X=4

The CEVLCCC tool will attempt to evaluate how cost-effectively each of the three payload designs can respond to changes in the threshold values of these five requirements (four existing, and one new). There is, of course, also the possibility that the program baseline will be changed in ways that cannot be reasonably foreseen at the present time. These so-called "unknown unknowns" are a genuine hazard for virtually every program; unfortunately, they are axiomatically beyond the scope of an a priori quantitative valuation strategy such as the CEVLCCC tool.

4.3. Marginal Probability Cost (MPC)

A key CEVLCCC assumption is that it is possible to formulate probabilistic modeling of the stochastic processes that cause deviations in the APB. One way to accomplish this is to treat the value for each performance parameter—in this case, each threshold value—as a random variable, and construct its cumulative distribution function (CDF). Then for each potential threshold value, there is an associated marginal probability within the CDF, as well as a corresponding LCC estimate to effect that capability for each design option. In aggregate, these cost and probability threshold descriptions comprise what we refer to as each requirement's *Marginal Probability Cost* (MPC).

Table 4 shows the MPCs for all three payload designs relative to requirement #1, air-to-air armament (for simplicity, the costs in these tables are shown as mean values rather than as a range of estimated values with an associated confidence interval that the actual CEVLCCC interface accommodates). The bolded row represents the current threshold value. We would generally not expect to have a dollar amount specified in this row for any design option, as any costs related to meeting the current threshold value would presumably have been incorporated into the traditional LCC estimates in Table 1. However, if there is uncertainty related to this threshold value, the standard LCC estimate may be adjusted accordingly, and the structure of the MPC matrix can accommodate that. Reductions in the requirement value (i.e., making the requirement less stringent) can also be accommodated.

Table 4. Marginal Probability Costs for Requirement #1 (A2A Armament)

Threshold Value (X)	Probability		Mean Estimated Cost ($M)		
	Cumulative	Marginal	Small Internal	Large Internal	External
8	5%	5%	$127.5	$46.0	$2.8
6	15%	10%	$72.5	$12.5	$2.3
4	**100%**	**85%**	**$0.0**	**$0.0**	**$2.0**

To illustrate how to interpret this table, consider the current threshold value of $X = 4$. The program has estimated that there is a 100 percent chance that the system will be required to have the capability to employ at least four air-to-air missiles (the current value). However, they have also estimated that there is a 15 percent chance the fighter will need to be able to carry at least six missiles instead (and a 10 percent chance it will need to carry *exactly* six). If such a requirement change occurs, there will be a cost impact regardless of the design chosen, but the level of impact varies greatly among the designs. If six missiles are required, the cost impact is relatively low for the *External* design ($2.3M) where there is ample space for expansion; moderate for the *Large Payload* design ($12.5M) where there is some extra space; and substantially greater for the *Small Payload* design ($72.5M) where there is no available space. In other words, the *External* payload design is the most flexible with respect to changes in the air-to-air armament requirement, and the *Small Internal* payload design is the least flexible.

It is important to recognize that the MPCs are time-dependent. Both the probabilities that a requirement will change and the costs incurred due to that change will certainly vary over time. Under a traditional acquisition strategy (as opposed to evolutionary acquisition), as a program matures, the probability that a requirement threshold value will change is likely to reduce, whereas the cost of accommodating it is likely to increase. The CEVLCCC tool is agnostic to the direction of probability change, but consider an example that might apply to the traditional acquisition model approach. A program might estimate that a particular threshold value has a ten percent cumulative probability of changing prior to the Preliminary Design Review (PDR), but only a five percent probability of changing between the PDR and Critical Design Review (CDR). Viewed in this way, the reader may recognize a certain similarity between these various MPCs and traditional risk burn-down plans. This is an important point, as the MPCs would need to be tracked in a similar manner, and could reasonably be integrated with traditional risk management techniques.

4.4. Constructing the CEVLCCC

Fundamentally, each MPC is an expected value calculation: however, since we intend to have the CEVLCCC output be a probability distribution, the intermediate expected values cannot simply be summed. Instead, we must track the mean and variance of all relevant constituent distributions as they are fused into the final curve that characterizes the stochastically-adjusted life cycle cost.

Given even a modest number of requirements, though, the potentially large number of associated threshold values along with the possibility of multiple time phases can quickly lead to a highly cumbersome set of distribution calculations. The current version of the CEVLCCC tool makes two simplifying assumptions to make this task more manageable. First, it assumes that the uncertainty associated with each cost estimate is normally distributed. Second, it treats the probability estimates associated with each threshold value as accurate point estimates with no associated uncertainty. Even allowing for these caveats, the calculations are still not trivial, so the remainder of this section will describe the underlying computational algorithm in some detail.

The first task is to merge the time points for the expected values of each requirement threshold value. To achieve this, we calculate the weighted mean and standard deviation values for each threshold value of each requirement across all time horizons. Next, we combine—or more formally, *convolve*—the expected value distributions for each requirement threshold into a consolidated requirement-specific MPC. Given the assumption that the expected values (i.e., the product of the costs and the probability) associated with each requirement threshold value are independent of one another, this becomes a relatively straightforward task. This is because we know that the convolution of independent, normally distributed probability density functions yields a normally distributed density function. Moreover, the mean and variance of the resulting density function are determined by summing the means and variances, respectively, of the original functions.

Once we have collapsed the expected values across time points into a single distribution and convolved all expected values within a given requirement into a single distribution, our last task is to convolve the expected values of each requirement into a single distribution that characterizes the LCC distribution of each design option. This last step yields a unique CDF for each design option, constructed from a normal probability distribution function with a known mean and variance. Once all the intermediate steps are accounted for, the comprehensive formula to obtain the mean—or expected value—of this distribution function is given as

$$E[CEVLCCC_d] = LCC + \sum_{i=1}^{R}\sum_{j=1}^{N}\sum_{t=0}^{P}\left(\left[(w_j)\left(E\left[C_j\right]\right)(p_j)\right]_t\right)_i \quad (1)$$

and the full standard deviation equation becomes

$$\sigma[CEVLCCC_d] = \sum_{i=1}^{R}\left[\sqrt{\left(\sum_{j=1}^{N}\sqrt{\sum_{t=0}^{P}\left(\sqrt{\left[(w_j)(\sigma_j)^2\right]_t}\right)^2}\right)^2}\right]_i \quad (2)$$

where

$C = cost\ range\ associated\ with\ given\ threshold\ requirement$
$p = probability\ associated\ with\ given\ threshold\ requirement$
$w = normalized\ weighting\ parameter$
$P = number\ of\ time\ points\ to\ be\ evaluated\ for\ a\ given\ requirement$
$N = number\ of\ threshold\ values\ to\ be\ evaluated\ in\ a\ given\ requirement$
$R = number\ of\ requirements\ to\ be\ evaluated\ (both\ existing\ and\ new)$
$d = design\ alternative$

From this, the CEVLCCC—in the form of a CDF—for each design alternative can be graphically represented and compared.

5. Discussion

5.1. Implementation

The CEVLCCC tool is intended for implementation at the individual program level. Logically, if this results in the selection of designs that have a greater value over the system's life cycle, the payoffs could be enormous. If the use of this tool resulted in selecting designs that were, on average, ten percent more cost effective over their life cycle, the Pentagon would easily save—or obtain value equivalent to—tens of billions of dollars every year.

While the potential benefits to the program are substantial, so too are the investment costs. There is undeniably a large amount of effort associated with obtaining valid cost and probability estimates for each potential threshold value associated with both existing and new requirements. Depending on the magnitude of the design decision, dozens of cost and probability estimates may need to be generated, each of which will consume project resources. Moreover, the utility of the CEVLCCC tool depends in large part on the first "C": It needs to remain current. In order to be able to make timely and relevant decisions, these numerous cost and probability estimates would need to be frequently updated, further burdening program resources.

In addition, regardless of the maturity or size of the program, we believe the up-front costs can be effectively managed in a number of ways. For instance, with respect to the cost estimates, while the program may occasionally prefer a formal cost proposal from the contractor, in most cases, the less rigorous (and far less costly) "engineering estimate" would likely be sufficient. In addition, it is reasonable to suppose there is significant synergy in the effort necessary to analyze the cost impacts of multiple threshold values relative to the same requirement. For the probability estimates, programs might derive values by maintaining awareness of emerging threats and technological breakthroughs, as well as studying historical trends of similar programs.

While the preceding points may bolster the case for implementing the CEVLCCC tool, the real key to making this type of approach feasible will require that defense programs adopt a *fundamentally different approach to risk management*. As noted earlier, there is a close relationship between the cost impact and probability estimates required for the CEVLCCC analysis and the information needed to perform traditional risk management. The salient difference is that risk management traditionally only considers the downside of uncertainty, whereas a stochastic LCC analysis must account for the upside of uncertainty as well, i.e., *opportunity*. Thus, the current practice of managing only the risk component of uncertainty is too narrowly conceived; programs should, in fact, have an *uncertainty* management plan that identifies sources of uncertainty, characterizes their ranges, and estimates probabilities of (both good and bad) outcomes [23]. Then the program can develop mitigation strategies related to both risks *and* opportunities. Under such an approach, not only would the types of inputs needed to maximize the utility of the CEVLCCC tool be more readily available, but the tool could be readily integrated into the program's overall uncertainty management strategy.

5.2. Future Work

The authors believe the principal improvement needed for the CEVLCCC tool is the ability to accommodate nonparametric distributions. This seems especially important for the cost estimates, for which the normal distribution assumption may be unwarranted. This is not only because every cost estimate must have a zero lower bound, but the uncertainty surrounding cost estimates tends to be asymmetric, such that the risk of exceeding the expected cost is generally not equal to the risk of a commensurate under-run. To implement this change would require additional information regarding each cost estimate, but even a minimal amount of input (e.g., specifying whether the distribution is positively or negatively skewed) would help to make the distributions more reflective of reality, and the final CEVLCCC more reliable. Although the convolution calculations required to accommodate non-parametric distributions would be cumbersome, the reliability improvement would likely be worth the effort.

Another key enhancement is to allow the user to characterize the uncertainties related to the threshold value CDF probabilities. This improvement would greatly complicate the underlying calculations, especially if the estimates were allowed to be non-parametric, but again the increased accuracy and reliability of the CEVLCCC may make the investment worthwhile. For both this improvement and the cost estimating improvement, integrating Monte Carlo analyses may be the most effective solution.

Finally, the CEVLCCC tool needs to be battle-tested. It should be validated via use in historical (vice hypothetical) case studies, and ultimately vetted through an active defense program in order to assess its practical efficacy. As part of this effort, we anticipate that users are likely to request the ability to more readily isolate the relative contributions of various inputs to the CEVLCCC (i.e., sensitivity analysis tools), and there are a number of additional features that could be embedded into the CEVLCCC tool to support that goal.

6. Conclusion

By assimilating and expanding the novel concept of LCC under uncertainty, the CEVLCCC tool presented in this paper is capable of serving as a straightforward, cost-based decision model for valuing system design options in the DoD. The authors have shown via a hypothetical use case how this tool can be used to quantitatively discriminate between designs using a stochastic version of expected LCC as a proxy for value. Under this approach, the best design is typically the one that is likely to be the most cost-effective over its life cycle.

Prior to introducing the CEVLCCC tool, we noted the problems related to using either NPV or real options techniques in defense applications, but we also asserted that the most formidable challenge to valuing flexibility in the DoD relates to monetizing military capabilities. The CEVLCCC approach sidesteps all of these issues. In fact, to the authors' knowledge, this tool represents the first quantitative methodology capable of justifying flexibility investments for Pentagon systems that does not need to assign value to the imputed capabilities or intended political outcomes. Moreover, the basic technique consists of a simple premise (i.e., expected value) and an intuitive output (i.e., life cycle cost), which can be readily understood by key stakeholders across the acquisition community, thereby potentially reducing entry barriers.

The CEVLCCC tool concept is premised on the notion that the need for capability changes in a program arises in a stochastic manner that can be modeled and incorporated into a continually updated, expected value model of total program cost. If implemented as part of an overall uncertainty management strategy, the authors contend that a tool like this could drastically improve design decisions in virtually all defense programs, and could feasibly reduce costs and/or improve value outcomes by tens of billions of dollars a year across the DoD.

References

1. Drezner J, Jarvaise J, Hess R, et al. An Analysis of Weapon System Cost Growth, Santa Monica, CA: RAND, 1993.
2. Christensen D, Searle D, Vickery C. The Impact of the Packard Commission's Recommendations on Reducing Cost Overruns on Defense Acquisition Contracts, *Acquisition Review Quarterly*, (Summer), 1999, 251-262.
3. Younossi O, Arena M, Leaonard R, et al. Is Weapon System Cost Growth Increasing?, Santa Monica, CA: RAND Corporation, 2007.
4. Rajan P, Van Wie M, Campbell M, et al. An Empirical Foundation for Product Flexibility, *Design Studies*, 26(4), 2005, 405-438.
5. Suh E, de Weck O, Chang D. Flexible Product Platforms: Framework and Case Study, 18(2), 2007, 67-89.
6. Shah N, Viscito L, Wilds J, et al. Quantifying Flexibility for Architecting Changeable Systems, in 6th Conference on Systems Engineering Research, Los Angeles, CA, 2008.
7. Saleh J, Mark G, Jordan N. Flexibility: A Multi-disciplinary Literature Review and a Research Agenda for Designing Flexible Engineering Systems, *Journal of Engineering Design*, 20(3), 2009, 307-323.
8. Kumar R. Understanding DSS Value: An Options Perspective, *Omega*, 27(3), 1999, 295-304.
9. Sivanthi T, Killat U. Valuing the Design Flexibility of a Distributed Real-time Embedded System, in International Symposium on Industrial Embedded Systems, La Grande Motte, France, 2008, 163-168.
10. Brown O, Eremenko P. Value-Centric Design Methodologies for Fractionated Spacecraft: Progress Summary From Phase 1 of the DARPA System F6 Program, in AIAA Space 2009 Conference and Exposition, Pasadena, CA, 2009.
11. Copeland T, Keenan P. How Much is Flexibility Worth, *McKinsey Quarterly*, (2), 1998, 38-49.
12. Ajah A, Herder P. Addressing Flexibility During Process and Infrastructure Systems Conceptual Design: Real Options Perspective, in IEEE International Conference on Systems, Man and Cybernetics, , 2005, 3711-3716.
13. Saleh J, Hastings D, Newman D. Flexibility in System Design and Implications for Aerospace Systems, *Acta Astronautica*, 53(12), 2003, 927-944.
14. Mayer Z, Kazakidis V. Decision Making in Flexible Mine Production System Design Using Real Options, *Journal of Construction Engineering and Management*, 133(2), 2007, 169-180.
15. Banerjee P, de Weck O. Flexibility Strategy--Valuing Flexible Product Options, in INCOSE/ICSE Conference on Synergy Between Systems Engineering and Project Management, Las Vegas, NV, 2004.
16. Ekstrom M, Bjornsson H. Valuing Flexibility in Architecture, Engineering, and Construction Information Technology Investments, *Journal of Construction Engineering & Management*, 131(4), 2005, 431-438.
17. Collopy P, Hollingsworth P. Value-Driven Design, in 9th AIAA Aviation Technology, Integration, and Operations Conference, Hilton Head, South Carolina, 2009.
18. Boden B, Ahlen A. Real Options Analysis: A Study of Implementation Impediments, Thesis. Gothenburg, Sweden: University of Gothenburg, School of Business, Economics, and Law, 2007.
19. Defense Acquisition University. Defense Acquisition Guidebook, [Online]. 2012. [cited 2012 February]. https://dag.dau.mil/Pages/Default.aspx.
20. Ryan E, Schubert C, Jacques D, Ritschel J. A Macro-Stochastic Model for Improving the Accuracy of DoD Life Cycle Cost Estimates, *Journal of Public Procurement*, 2013,.
21. Brown O, Long A, Shah N, Eremenko P. System Lifecycle Cost Under Uncertainty as a Design Metric Encompassing the Value of Architectural Flexibility, in AIAA Space Conference, 2007, 216-229.
22. Brown O, Eremenko P. Application of Value-Centric Design to Space Architectures: The Case of Fractionated Spacecraft, Reston, VA: American Institute of Aeronautics and Astronautics, Report No.: AIAA-2008-7869, 2008.
23. de Neufville R, Scholtes S. Flexibility in Engineering Design Cambridge, MA, MIT Press, 2011.