SYM-AM-17-059

Proceedings of the Fourteenth Annual Acquisition Research Symposium

Wednesday Sessions
Volume I

Acquisition Research:
Creating Synergy for Informed Change

April 26–27, 2017

Published March 31, 2017

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

Acquisition Research Program
Graduate School of Business & Public Policy
Naval Postgraduate School

Determining the Value of a Prototype

Zachary S. McGregor-Dorsey—received his degree in mathematics from the University of Colorado Boulder. He has worked in the DoD cost-related fields for nearly a decade, working on everything from research to detailed life cycle cost estimates. His current interests lie in graph theory, machine learning, and banjo playing, and he is always looking for applications of those interests to the cost world. [zmcgrego@ida.org]

Abstract

Most major defense acquisitions require a technology that is not yet fully developed, introducing a non-negligible amount of risk to the program and its cost. Developing a prototype or technology demonstrator prior to execution of the program can be useful in mitigating this risk, yet these demonstration programs also have associated costs. This paper develops a method to value this risk mitigation, setting an appropriate maximum cost for the demonstration program. This novel application of Value of Information theory and properties of the Bayesian preposterior distribution requires only a program cost estimate distribution and some estimate of possible results of the demonstration program. The method is broadly applicable to programming with varying amounts of technological uncertainty. We describe the method, then show how actual cost overruns of historical programs with and without prototypes can be used to estimate a value of prototype efforts relative to estimated program cost. We conclude with a discussion of other applications and how to explain the method and results to decision-makers.

Introduction

With any new defense program, the decision to begin its research, development, test, and evaluation (RDT&E) phase is a significant commitment of resources. It is a commitment that—due to contracts, regulations, and momentum—is not easily annulled. Even programs that experience cancellation tend to do so after appreciable investment. The decision-maker faces this weighty determination at a time when information about the program and its challenges is largely speculative and difficult to measure, making any estimate of its cost subject to a high degree of uncertainty.

In this situation, a decision-maker may appreciate additional information that helps to refine their estimate of the resources necessary to complete the program. A possible source of this information is a prototype or technology demonstrator.[1] A key attribute of a prototype program is its possible existence outside a formal RDT&E program; the information gained can then be used to make the riskier decision on moving ahead with the program.[2] However, prototype development itself can be costly, so the decision-maker would prefer to know in advance the value of the information that would be provided by the prototype effort.

This paper is intended to address this wish, providing a method to value the prototype effort using the decrease in uncertainty it would provide to the cost estimate. To employ the method, the decision-maker needs only an initial cost estimate (presented as a

---

[1] Hereafter, we refer to both prototypes and technology demonstrators as simply *prototypes* because we are only interested in their role as information augmenter. Indeed, any program or project that can reduce the uncertainty of a cost estimate is a good candidate for the method herein.

[2] Even if the prototype occurs as part of RDT&E, it still may offer a natural point for reassessment, depending on the general shape of the program.

distribution of possible costs), a cost constraint for the program, and some estimate of the decrease in uncertainty the prototype effort would provide. We introduce the method through increasingly complex illustrations of its application, first employing Value of Information (VoI) theory to establish a method for comparing the prototype/no prototype options, and then using properties of the *preposterior distribution* (defined below) to compare the values. We end the paper by using the method to create a rule of thumb, developed from historical data, for the value of a typical prototype.

Illustrations of Method

In this paper, we are concerned with the particular application to determining prototype value in the case of defense acquisition. We use language that reflects this focus, even though much of the work below is more broadly applicable. We also note that this paper presents a nontechnical description of the method; some detail and mathematics are omitted and can be found in the references.

Framework and Terminology

Before a defense acquisition program begins, a decision-maker is presented some estimate of the final cost[3] of the program. This estimate will be a probability distribution of possible final costs; we call this the prior cost distribution and denote the final cost subject to it as the random variable *C*. The decision-maker also has some cost constraint *b* for the overall program, presumably related to the relative importance of the associated mission and the priority of the program within its portfolio. The cost constraint *b* can be interpreted in several ways. It might simply be the maximum amount the decision-maker is willing to spend to address the mission. More realistically, it might be the known cost of a lower-risk alternative that addresses the same mission need, or the point at which the opportunity cost of choosing the program over a lower-risk, less capable system is greater than the capabilities gained.[4] For simplicity, we assume the cost constraint to be the cost of satisfying the mission the program addresses in some other way.

The cost constraint informs the decision directly: If the final cost of the program is expected to exceed *b*, i.e., $C > b$, the decision-maker should not embark on the program; otherwise, when $C \leq b$, they should. Of course, the final cost *C* has some probability of being greater than *b* and some probability of being less, but the decision-maker cannot make a fraction of a decision. For simplicity, we assume the decision-maker relies on the expected value[5] of $C, \mathrm{E}[C]$, to make their decision.[6] That is, the program will only be considered if $\mathrm{E}[C] \leq b$, as illustrated in Figure 1.

---

[3] While by no means necessary, it is helpful to think of this final cost as the cost of the associated RDT&E program, because the relationship between RDT&E and a prototype is more direct.
[4] As may be inferred from the interpretations, it is possible that the cost constraint itself is a probability distribution of costs. This is a possible extension and we will discuss letting *b* vary. However, for simplicity, we assume the cost constraint is a single value.
[5] Recall that the expected value of a random variable is the average value resulting from (infinite) repeated sampling of the random variable. This is calculated as the sum of all possible values of the random variable weighted by their associated probabilities.
[6] Other constraints may be used, such as $Pr(C \leq b)$ for some probability p. However, as will be seen, to do so requires assumptions about the shape of the resulting cost distributions that the expected value assumption does not.

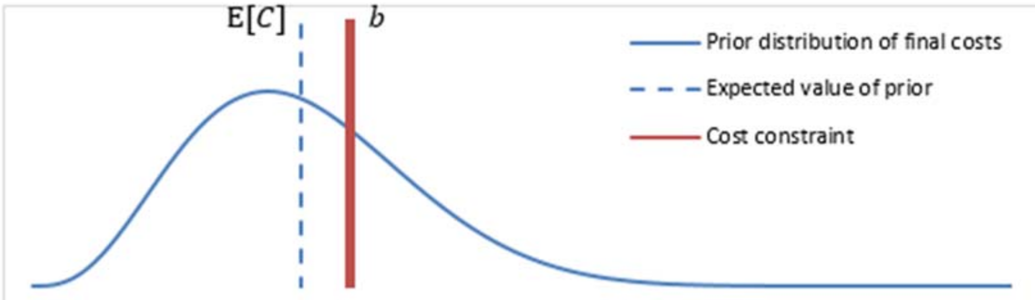

Figure 1. Example of a Program That Should Proceed

We introduce the option of obtaining more information about the distribution of final costs through prototype development. The decision-maker now has three choices:[7]

1. Start the program (with full intention to complete it);
2. Pursue the alternative program at cost *b*; or
3. Invest in a prototype, after which another decision will be made to start or not start the program.

The decision-maker can choose between the first two options using the cost constraint as discriminator. However, the closer in value E[*C*] and *b* are, the greater the probability that *b* lies between E[*C*] and the actual observed final cost, meaning the decision was not optimal. As these values approach each other, the decision-maker then loses confidence in their decision and becomes more interested in obtaining more information. At some point, they instead choose the third option.

The third option will, at some investment cost *I* of the prototype program, provide the decision-maker with a new probability distribution of possible final program costs. We call this a *posterior cost distribution* and denote the final (uncertain) program cost subject to it by the random variable *D*, adding a subscript when more than one posterior cost distribution is referenced. We assume *D* is the remaining cost of the program after completing the prototype, i.e., the cost *I* is not reflected in *D*. After a prototype program is pursued, the decision-maker will make the decision to continue the main program similarly to the one described previously, but now comparing E[*D*] and $b - I$.

Note that before the prototype is developed, its resulting *D* is unknown and, indeed, many different posterior cost distributions D_i might arise. Each of these distributions has some probability of occurring, so we can consider the probability distribution over posterior cost distributions, to which we associate the random variable Δ. Abusing notation a bit, we can say the values of Δ are all the possible random variables D_i. We will actually be more interested in the probability distribution of the expected values of $\mathrm{E}[D_i]$ of posterior cost distributions. We associate the random variable $\bar{\Delta}$ with this distribution of means.

The decision-maker is interested in pursuing a prototype only when the result is likely to change their decision. If their decision is not changed, they stay on the same path, as

[7] Of course, specific real world situations are more complicated, requiring some adjustment of the model.

initially they gain nothing more than confirmation and the cost *I* of the prototype is wasted.[8] For example, consider Figure 2. While the decision-maker does not know what the posterior cost distribution will be, they do know it is likely to have an expected value somewhere in the thick part of the prior cost distribution (blue line). Otherwise, we could not consider the prior cost distribution to be an accurate reflection of our prior knowledge, and should change it accordingly. The example posterior cost distribution (orange line) given in the figure is plausible, though less likely than one in the thicker part of the distribution.

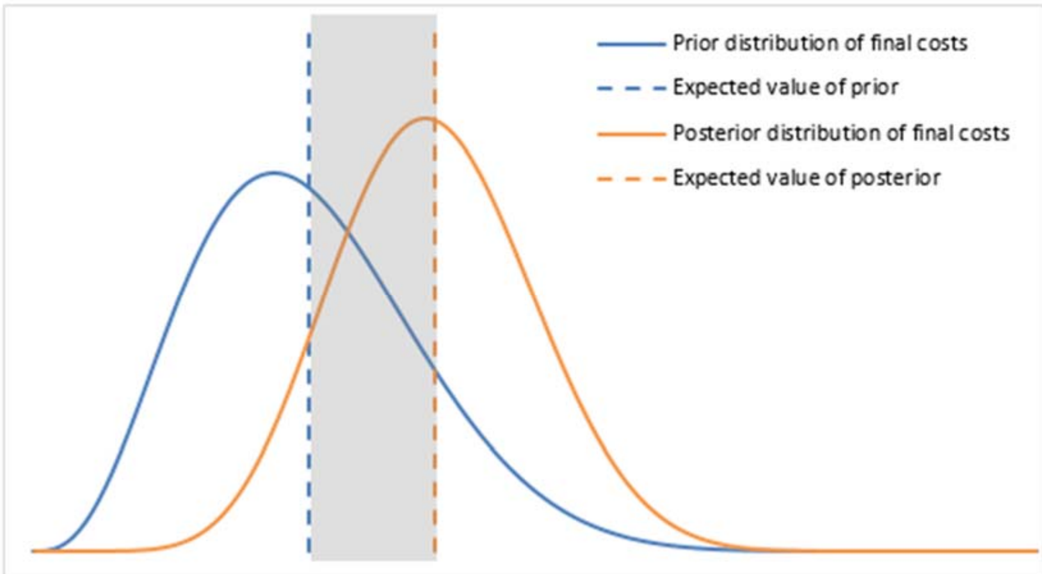

Note. This decision changes with the added information of the posterior distribution whenever the cost constraint *b* lies in the shaded region.

Figure 2. When the Posterior Cost Distribution Is Useful

Suppose the constraint *b* is very high, far to the right in the chart. Here, the decision-maker would start the program, almost without question. If the initial distribution *C* is accurate, it would be very unlikely that a prototype program would reveal a new expected value above that constraint, so a prototype is unlikely to change the decision. In this case, the value of a prototype is very low. Similarly, if *b* is very low, far to the left in the chart, the prototype is very unlikely to change the decision to not start the program. In both cases, for the specific case of *D* shown in the chart, the value of the prototype is 0.

On the other hand, when the constraint *b* more evenly splits the prior cost distribution, it is more likely that a prototype could change the *a priori* decision. In Figure 2, if *b* took any value in the shaded region, the decision would change after that particular prototype result. We can quantify the expected value of the prototype effort using the difference in the expected costs to the decision-maker between making a prototype or not.

[8] It is likely that the program will benefit from prototype spending if the program is continued, so the entire cost of the prototype may not be lost, but we do not address this directly. The benefit from the prototype is assumed to be captured in the increased fidelity of the cost distribution.

For this calculation to be intuitive, we need to look at all possible posterior cost distributions together, using VoI theory.

Doing so gives us a distribution of values for the prototype determined by Δ, *C*, and the constraint *b*. We call the expected value of this distribution the value of the prototype. We will assume *C* and *b* are given, and use the *preposterior distribution* of program costs together with assumptions about the amount of information a prototype provides, in order to determine this value.

We make the following assumptions:

- Decisions regarding the continuation of the program are based on expected values.
- The distribution of final costs *C* is reliable, i.e., unbiased with unbiased estimation of standard deviation.
- Any uncertainty in the cost of the prototype program is small compared to the cost of the entire program, and thus can be treated as a point value.
- A firm affordability cap for the program is known and expressible as a constant bound.
- Modifications can be made to the method to relax these requirements, but this is not addressed here.

One assumption requires specific attention. The requirement that the distribution of final costs is known and accurate is strong. For example, this implies the program is technically feasible and will, if attempted, succeed. Additionally, there is historical evidence that initial cost estimates are systematically low. (See the section titled A General Application for some evidence of this.) Even if the expected value of the estimate is correct, the variance (risk) is not all known and can therefore be vastly underestimated. We have not tested the robustness of this method against inaccurate initial cost estimates, although we hope to in the future. For that reason, any application requires great care and the most conservative of initial cost estimates.

Value of Information (VoI)

VoI is a concept from decision analysis. In its simplest form, the theory examines two possibilities, one in which a decision is made without a certain piece of information and one in which it is made with the information. The decisions have costs and benefits, depending on some unknown state of the future, which the information describes in some way. The difference between the expected values of these decisions is the value of the information.

A Simple Illustration

For illustration, suppose the decision-maker has the information that a program has a 40% chance of costing \$9 billion (\$9B), and a 60% chance of costing \$14 billion (\$14B). The expected cost of the program is thus $0.4 * \$9\text{B} + 0.6 * \$14\text{B} = \$12\text{B}$. If their cost constraint *b* is \$10 billion, they will not initiate with the program as it stands.

However, they are then told that a prototype effort would give them the information they need to determine whether the full program is a \$9 billion or \$14 billion program. How much should they be willing to spend on the prototype?

Recall that the cost constraint of \$10 billion is the value of satisfying the mission, which, for purposes of explication, we assume to be the cost of satisfying the mission in some other way. If the prototype is not pursued, the decision-maker will pursue the \$10 billion alternative, and so the expected cost of meeting the mission need is \$10 billion.

If the prototype is pursued, there are two possibilities. The prototype will reveal—40% of the time—that the program will cost $9 billion. The decision-maker will then spend $9 billion to execute the program. The remaining 60% of the time, the prototype will reveal a cost of $14 billion. The decision-maker will then stick with the $10 billion alternative. Therefore, when the prototype is pursued, the expected cost of satisfying the mission is $0.4 * \$9\text{B} + 0.6 * \$10\text{B} = \$9.6\text{B}$.

By VoI, we conclude the value of the prototype is $\$10\text{B} - \$9.6\text{B} = \$0.4\text{B}$ to the decision-maker, the difference between the two decision paths. That is, the decision-maker should be willing to spend up to $0.4 billion dollars on a prototype that can provide this level of information.

This illustration, however, is far from realistic. The distribution of final costs *C* is not, as a rule, discrete, and the prototype effort will typically not provide perfect information on the final cost. These observations can be accommodated with a small amount of generalization, as we will see in the next illustration. However, we will assume for now that we know what the distributions in Δ look like. This assumption will be addressed in the next section.

Illustration of More Complexity

Now suppose the decision-maker is presented with any distribution of final costs *C* such that E[*C*]=$11 billion. Suppose further that the prototype will demonstrate two critical technologies, A and B, and each technology will be determined to be easy or difficult to mature. Therefore, from the prototype, four new cost distributions can result. Letting the random value indicate its distribution, we get $\Delta = \{ D_1, D_2, D_3, D_4\}$, where

D_1 corresponds to both technologies proving easy to mature,

D_2 corresponds to A being easy and B being difficult,

D_3 corresponds to A being difficult and B being easy, and

D_4 corresponds to both being difficult.

Prior to the prototype development, the conditional expected cost under each case was determined, and the probabilities of A and B being easy to mature are found to be 0.6 and 0.1, respectively, resulting in Table 1.

Table 1. Properties of the Posterior Cost Distributions in an Illustration of More Complexity

Distribution	Probability of Distribution	$\text{E}[C\|D_i]$
D_1	0.06	$8 billion
D_2	0.54	$9 billion
D_3	0.04	$13 billion
D_4	0.36	$14 billion

Given the cost constraint *b* of $10 billion, the decision-maker can now compute the value of the prototype. If they decide to not pursue the prototype, the mission cost is $10 billion, as before. With the prototype, note that they will pursue the program if distributions D_1 or D_2result, and cancel the program if distributions D_3 or D_4 result. The expected mission cost given a prototype effort is thus

$$0.6 * \$8B + 0.54 * \$9B + 0.04 * \$10B + 0.36 * \$10B = \$9.34B.$$

Therefore, the value of the information from the prototype is, in billions of dollars,

$$\$10B - \$9.34B = \$0.66B.$$

We note here that it is straightforward to apply sensitivity analysis to the cost constraint. In Figure 3, we see what happens as the cost constraint varies. The closer the constraint is to the $\mathrm{E}[C]$of \$11 billion, the greater the value of the prototype, because the constraint is more likely to be between $\mathrm{E}[C]$ and $\mathrm{E}[D]$, for any given posterior cost distribution of final costs *D*. We also see that as the cost constraint approaches the extremes, the prototype has no value, because the decision-maker would not change their decision from that based on *C* alone. The values in this figure together with the distribution of the cost constraint give the expected value of the prototype.

The above illustration is instructive, but it is unrealistic in one important way. In general, we will know what the distributions in Δ look like or how they are distributed, i.e., what their various probabilities of occurrence are. However, given our assumptions about how the decision will be made, we really only need to know the distribution of the *means* of the distributions in Δ. The next section will illustrate this.

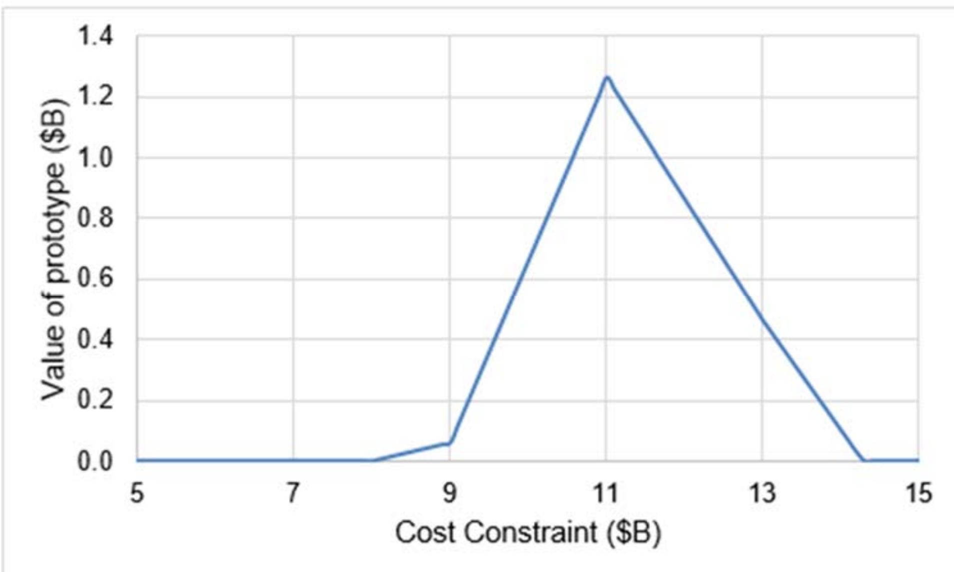

Figure 3. Value of Prototype as the Cost Constraint Varies

Preposterior Distribution

Our main goal in this section is to describe the distribution of means of posterior cost distributions $\bar{\Delta}$ which we call the *preposterior* cost distribution. (The initial distribution of final costs *C* is the *prior* cost distribution and a posterior cost distribution of final costs *D*, when known, is the *posterior* cost distribution.) Using this notation in application to an illustration of more complexity (discussed in the previous section), we have, in billions of dollars,

$$E(\bar{\Delta}) = 0.6 * \$8B + 0.54 * \$9B + 0.04 * \$13B + 0.36 * \$14.3B = \$11B.$$

Before we present the properties technically, we examine them intuitively. This section is largely derived from Section 5.4 of Raiffa and Schlaifer (1961); any additional justifications may be found there.

Creating a prototype provides information about the technical hardship of the program, which should have some effect on the prior cost distribution of *C* of that program. There is no reason, *a priori*, to expect this new information will increase or decrease the mean of the initial estimate; if there were, the initial estimate should be adjusted accordingly. We emphasize here that a specific posterior cost estimate might have a different mean from the prior cost distribution, but the mean of all possible posterior estimates must match that of the prior estimate if it is reliable. That is, the expected value of $\bar{\Delta}$ is the same as that of *C*, because *C* summarizes all of the information we have about what $\bar{\Delta}$ could look like.

The new information gleaned from the prototype will also produce a tighter cost distribution (lower variance) than that of the prior cost distribution.[9] Let us consider the extreme possibilities to gain intuition for this tightening. If the prototype provides no new information at all, any value of Δ is just *C*—the posterior is the same as the prior. So $\bar{\Delta}$ would consist of a single value, E[*C*], and clearly E $[\bar{\Delta}]$=E[*C*]. Thus, the variance of $\bar{\Delta}$ would be 0 and the variance of every possible value Δ takes on (every *D*) is equal to the variance of *C*.

On the other hand, if the prototype provides perfect information (i.e., tells us the exact cost of the program), the prototype gives the same information as actually doing the program. Any value *D* of Δ is the final cost. It is a degenerate distribution with all mass at that final cost, and these values must be random draws of *C*, since the prior cost distribution is accurate. The expected values of these *D* are the final cost, so $\bar{\Delta}$ consists of expected final costs taken from the prior cost distribution, i.e., $\bar{\Delta}$=*C*. Clearly, then, E $[\bar{\Delta}]$=E[*C*] and the variance of $\bar{\Delta}$ and *C* are equal. The variance of the degenerate distributions Δ definitionally have variance 0.

Taken together, we see that as the variance of the resulting distributions decreases (i.e., as the prototype provides more information), the variance of $\bar{\Delta}$ increases and is bounded by that of the variance of *C*. Indeed, the more information provided, the more $\bar{\Delta}$ looks like *C* as a distribution.

The observations above can be summarized with the following three theorems. Stating these theorems exactly requires some technical complexity beyond the scope of this paper. See Raiffa and Schlaifer (1961) for proofs.

> *Theorem 1. The expected value of the means of new cost distributions resulting from the information gained by a prototype is equal to the expected value of the prior cost distribution, i.e.,*
>
> $$E(\bar{\Delta}) = E(C)$$
>
> *Theorem 2. The variance of the means of new cost distributions resulting from the information gained by a prototype is equal to the variance of the prior cost distribution less the expected value of the variances of the new cost distributions, i.e.,*

---

[9] Intuitively, this makes sense; the more the decision-maker knows, the more specific they can make their estimate. However, this fact actually relies on the prior cost distribution being accurate, meaning no new information can be added that adds new possibilities. If something has not been thought of, this is captured by a wider variance and new information will only discount some of those possible outcomes.

$$Var(\bar{\Delta}) = Var(C) - E(Var(D))$$

Theorem 3. As the information provided by the prototype approaches perfect information, the distribution of $\bar{\Delta}$ approaches the distribution of C, i.e.,

$$\lim_{n\to\infty} Pr(\bar{\Delta} < d) = Pr(C < d)$$

where n is the amount of information.

The third theorem depends on our assumption that *C* is a reliable distribution, unbiased in mean and standard deviation. It also assumes that we can always add information that will sharpen our estimate. The second assumption is true because we can always run the experiment of completing the program, at which point the cost will be known.

The open problem at the end of the previous section was the shape of the distribution Δ, without which our VoI calculation would be impossible. While we are not able to describe this distribution completely, we now know some key attributes that can help us estimate the value of the prototype.

Putting It Together

For our final illustration, assume the decision-maker has a particular goal in mind for their prototype. Suppose they have a cost estimating relationship (CER) based, in part, on a particular technical attribute that is not precisely known, and their prototype will determine this attribute nearly perfectly.[10] Using the CER, they use Monte Carlo simulation with multiple possible values of the technical attribute (and the standard error of the CER) to estimate the cost distribution *C* of their program. For this example, assume that *C* has a scaled beta distribution with support on (0, $40B), mean $11B, and variance 40.[11] They then apply the CER to specific values of the technical attribute and find that the variance of the posterior, on average, is 10.[12]

From Theorem 1, the decision-maker knows the expected value of $\bar{\Delta}$ to be $11 billion. From Theorem 2, they can compute the variance of $\bar{\Delta}$ to be 30, the difference of the variance of *C* and the mean of variances of elements of $\bar{\Delta}$. Finally, from Theorem 3, lacking other information, they assume that $\bar{\Delta}$ is the same type of distribution as *C.* They can now estimate the value of the prototype.

Again, we assume the cost constraint is $10 billion. Without the prototype, they do not proceed with the program, so the cost of satisfying the mission is $10 billion. With the prototype, they consider two situations. In the first, the prototype reveals a cost less than the constraint. They then sum these costs, weighted by the probability they occur (given by the distribution of $\bar{\Delta}$), i.e., they calculate

$$\int_0^{10} x \, \mathrm{f(x)} dx = 3.1,$$

---

[10] For example, models may predict a hypothetical aircraft design has a lift-to-drag ratio in a certain range, but the exact ratio is not known. A prototype may demonstrate the actual lift-to-drag ratio.
[11] The correct units for variances here is squared billion dollars. We omit the units for clarity.
[12] In this situation, the decision-maker could just generate some representative subsample of Δ and apply VoI directly. We have simplified this example from one in which the reduction in variance after the prototype can be estimated but the distributions are not known, for purposes of illustration.

where $f(x)$ is the probability density function of Δ. In the second, the prototype reveals a cost greater than the constraint, so in each case, the cost of satisfying the mission is $10 billion. Thus, they calculate

$$\int_{10}^{\infty} 10\, \mathrm{f(x)}dx = 5.2.$$

The sum of these is the expected cost of satisfying the mission after the prototype. In our case, the expected cost is $8.3 billion. Thus, the decision-maker should be willing to pay up to $1.7 billion for the prototype.

We again consider what happens when we vary the cost constraint in Figure 4. As expected, the value of the prototype is greatest when the constraint is near the expected value of the prior cost distribution.

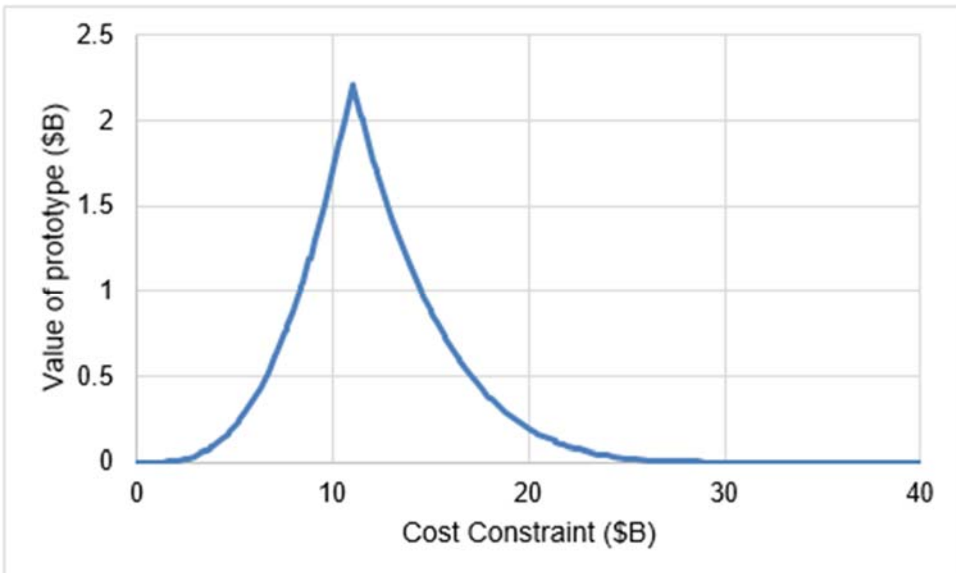

Figure 4. Value of Prototype as the Cost Constraint Varies

A General Application

Usually, the decision-maker does not have access to the decrease in variance provided by a prototype, so it is useful to develop a rule of thumb for the general case. To do so, we use historical cost overruns of development programs[13] with and without prototypes to estimate the change in variance. We take our data from Tyson, Nelson, Gogerty, Harmon, and Salerno (1991), which examined the development costs of 51 historical aircraft and tactical munition programs, 35 of which did not have prototypes and 16 of which did. The report found that, on average, final development costs for programs without prototypes was 1.62 ± 0.96 the cost of their initial cost estimate, whereas programs with prototypes were 1.17 ± 0.17 their initial cost estimate.

[13] We look at development programs here, not the total program, mainly due to data limitations. The effect of the prototype program on production and sustainment also needs consideration.

While the cost overruns are important, our goal is simply to compare the variances to find the value of a prototype. Suppose we have a new (possible) program. We treat the set of programs without prototypes, their costs scaled to their initial estimates to make comparable, as describing its initial cost estimate distribution of *C.* In doing this, we are treating these programs as observations of possible outcomes of our hypothetical program. Note this does not imply we expect this program to overrun by 62%; we are using this database as a reasonable proxy for an accurate prior cost distribution. We can choose our own expected value, so for simplicity, we assume E[*C*] is 1 of some unit. Scaling the variance accordingly, we find this distribution has a variance of 0.34.

We then treat the set of programs with prototypes as a sample of the set $\bar{\Delta}$ of possible means of distributions resulting from a prototype. Because $\bar{\Delta}$ has the same mean as *C*, we normalize these data to 1 as well and find $\bar{\Delta}$ has a variance of 0.02.

Using the above, we can now estimate the value of a prototype for different cost constraints. We assume a beta distribution with expected value of 1 and respective variances for *C* and $\bar{\Delta}$. (We note that the distribution of the data is well approximated by the beta distribution.) In Figure 5, we present the value assuming the expected value of *C* is 1. In Figure 6, we present the value of the prototype as a percent of cost constraint. Note that while the figures are similar, Figure 6 shows that the value of the prototype is relatively higher for smaller cost constraints.

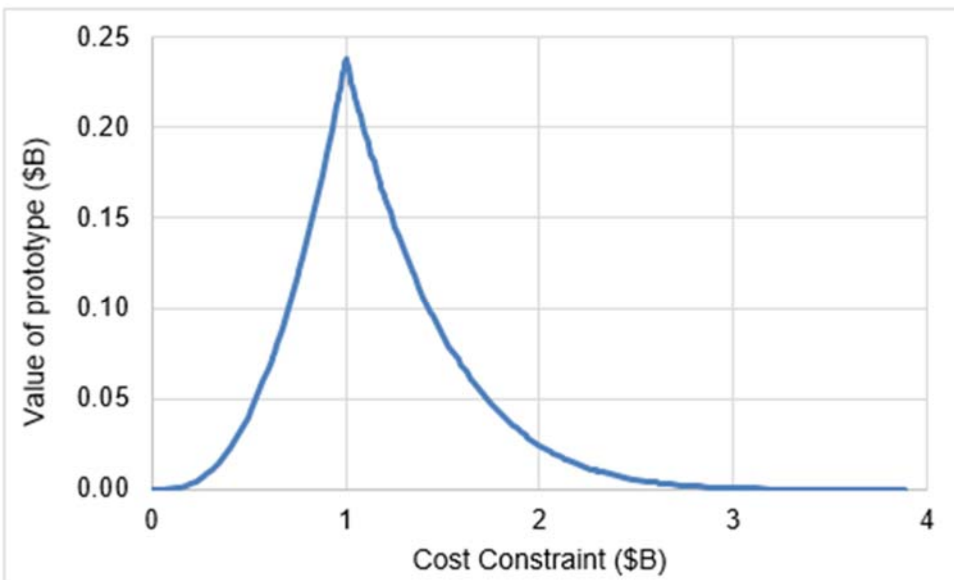

Figure 5. The Value of Prototype for a $1 Billion Development Program, Using Historical Development Cost Distributions

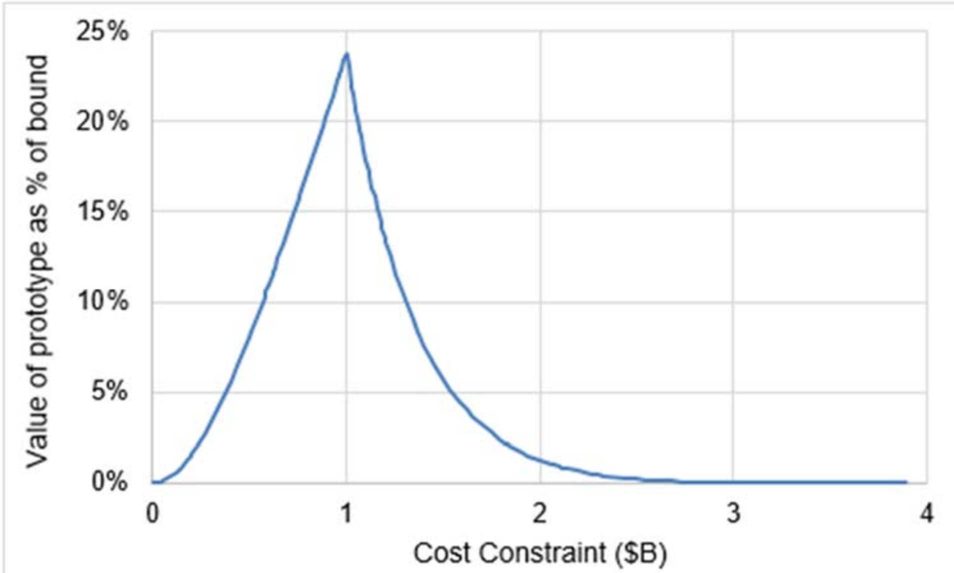

Figure 6. The Value of the Prototype Relative to the Cost Constraint, Using Historical Cost Distributions

A possible problem with the above formulation is that we only consider programs that successfully start and do not experience cancellation. The underlying data can be improved by adding failed programs and should be addressed in future research. Also, while it is safe to assume that none of the programs in the historical data benefitted from this research, it is unlikely that the decision of whether or not to build a prototype was random; that decision process could introduce a bias. On the other hand, it is possible that the decision on building a prototype was primarily determined by the political situation at the time and may therefore be entirely independent of the technical merits, in which case the sample would be unbiased.

Conclusion

This work needs expansion in several directions. Our immediate interest is in determining the effect of a biased prior. That is, assuming the distribution of C is unbiased in mean and standard deviation is a very strong assumption. When that assumption is violated, how much is the final value of the prototype affected? To this end, we also want to incorporate more historical data, particularly the actual cost of prototype programs, to test and refine the method in less than perfect cases.

We also would like to allow the decision-maker to use values other than the expected value for decisions. We do not see this as impossible, but it may be technically difficult and rely even more on the assumption that the distribution of $\bar{\Delta}$ is similar in shape to the distribution of C. Other directions of research should allow more generality in the method, such as allowing the cost constraint to also come from a distribution. Specific application to commodity type should also be studied to find likely reductions in cost variance due to prototype development. As always, additional historical data would also be a benefit.

The method presented above is a relatively simple way of determining the value of a prototype. The assumptions, while strong, are not unreasonable. In general, if we cannot assume that an initial cost estimate distribution is reliable, we should adjust the estimate to reflect our uncertainties. Estimating the amount of variance reduction provided by a prototype program is likely possible in cases where the prototype is geared to answering specific technical questions, which is most common. Indeed, the effect on the variance may provide a guide to what goals the prototype should have. In the cases where this variance reduction is not estimable, we can use historical programs to estimate the variance reduction. As shown above, the value of a prototype can reach more than 20% of the value of the entire program in those cases where the affordability of the program is most uncertain.

The author has developed several models in the course of this work, using both Excel and Python, that are available upon request.

References

Raiffa, H., & Schlaifer, R. (1961). *Applied statistical decision theory*. Cambridge, MA: MIT Press.

Tyson, K. W., Nelson, J. R., Gogerty, D. C., Harmon, B. R., & Salerno, A. W. (1991). *Prototyping defense systems* (IDA Document D-1097). Alexandria, VA: Institute for Defense Analyses.

Acknowledgments

The research for this paper was funded by the Air Force Research Laboratories. We thank them for their support.

Acquisition Research Program
Graduate School of Business & Public Policy
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net