INTERVAL CHARACTERIZATION OF SELECTED EPISTEMIC UNCERTAINTIES

Chris Everett[1], Homayoon Dezfuli[2], Bob Youngblood[3], Tony Hall[1]

[1] *ISL, Rockville, MD, USA*
[2] *NASA, Washington, DC, USA*
[3] *INL, Idaho Falls, ID, USA*

The conventional method for treating epistemic uncertainty within a probabilistic model is to assign probability distributions to the uncertain parameters that express relative belief about the true state of the world. However, there are objections to the use of probability distributions in this way, including arguments that there are situations where there is not enough information to choose a distribution, and that a formalism that requires a distribution has the undesirable characteristic of forcing unwarranted assumptions into the analysis. This study explores some of the implications of characterizing some epistemic uncertainties as intervals, rather than as distributions, within an otherwise distribution-based model. It focuses on implications for decision making, looking into what an interval characterization means in terms of meaningful decision rules that can be brought to bear when selecting from different options.

The probabilistic model used in this analysis is the ascent phase of a space launch vehicle with abort capability, where the mean abort lead time is epistemically uncertain over a specified interval. The decision used for this study is that of integrated vehicle health monitoring system selection, where the choice is between two options: Option A, which has a relatively low false positive error rate but is believed to provide relatively short abort lead times, leading to lower abort effectiveness given launch vehicle failure; and Option B, which has a relatively high false positive error rate but is believed to provide longer abort lead times and higher abort effectiveness.

The interval characterization of mean abort lead time propagates into the analysis, resulting in interval characterizations of key quantities of interest to decision makers, such as expected utility and value of perfect information. In cases where intervals overlap, rational decision making is thwarted by the inability to unambiguously rank options, i.e., in terms of their expected utilities.

Intervals used to characterize parameter uncertainty propagate into the statistical quantities that underpin rational decision making, necessitating the use of alternative decision rules. One such decision rule, which is consistent with a risk-averse attitude, is to select the option for which the minimum possible expected utility is highest. Such a maximin decision rule focuses on minimizing the potential downside of a decision, without regard for the magnitude of the potential upside. The maximin rule may be most appropriate when striving to achieve some threshold utility, such as when some threshold level of performance is required of a system, so long as the minimum performance exceeds the required level. Other decision rules are possible, and can lead to different decisions. A decision rule that reflects a risk-seeking attitude might be to select the option with the highest maximum possible expected utility.

I. INTRODUCTION

The conventional method for treating epistemic uncertainty within a probabilistic model is to assign probability distributions to the uncertain parameters that express relative belief about the true state of the world. However, there are objections to the use of probability distributions in this way, including arguments that there are situations where there is not enough information to choose a distribution, and that a formalism that requires a distribution has the undesirable characteristic of forcing unwarranted assumptions into the analysis. One such situation is the choice of a prior distribution to express belief about something for which there is not much hard evidence, in which case the analysis can potentially be sensitive to highly subjective opinions that can vary from one individual to another. The use of so-called non-informative priors is often used to minimize the introduction of unwarranted assumptions, but even so there are issues, as discussed in Guarro[1].

Ferson[2] discusses a number of methods for addressing epistemic uncertainty in modeling, beyond the use of probability distributions. In particular, he discusses methods that involve characterizing epistemic uncertainty using intervals, in which statements are made concerning the minimum and maximum credible values of an epistemically uncertain parameter, but nothing is said about the relative probability that the true value of the parameter is at any particular location within the interval.

This study explores some of the implications of characterizing some epistemic uncertainties as intervals within an otherwise distribution-based model. It focuses on implications for decision making, looking into what an interval characterization means in terms of meaningful decision rules that can be brought to bear to select among different options. It also takes an initial look at how the concept of value of information (VOI) might be applied to a model that contains interval-characterized epistemic uncertainties.

II. THE MODEL

The probabilistic model used in this analysis is the ascent phase of a space launch vehicle (LV) with abort capability, where there is the possibility of LV failure resulting in loss of mission (LOM) if the abort is successful and loss of crew (LOC) if the abort fails.* Additionally, the abort system might mistakenly register LV failure (i.e., a false positive), in which case an abort is initiated from an otherwise functional LV, leading to LOM if the abort is successful and LOC if it is not. Figure 1 illustrates the situation in event tree form, showing the various paths leading to the three possible end states of LOC, LOM, and OK.

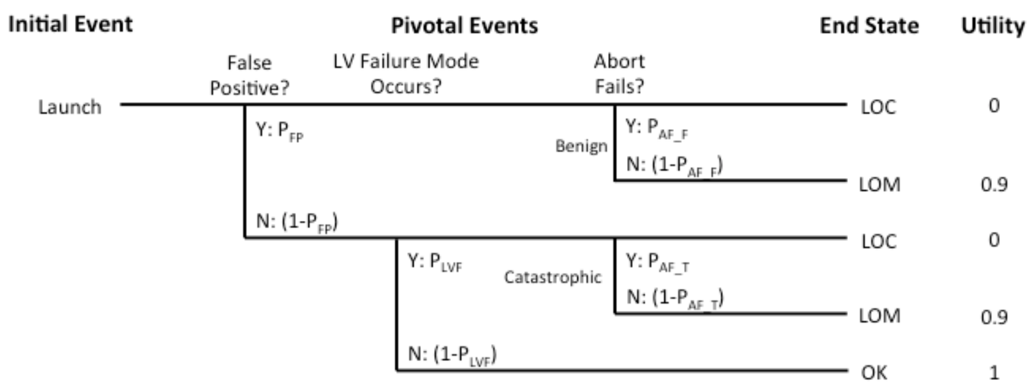

Fig. 1. LV Ascent Event Tree

The probabilities of LOC, LOM, and OK are:

$$P(LOC) = P_{FP} * P_{AF\_F} + (1 - P_{FP}) * P_{LVF} * P_{AF\_T} \tag{1}$$

$$P(LOM) = P_{FP} * (1 - P_{AF\_F}) + (1 - P_{FP}) * P_{LVF} * (1 - P_{AF\_T}) \tag{2}$$

$$P(OK) = (1 - P_{FP}) * (1 - P_{LVF}) \tag{3}$$

where:

P_{LVF} = Probability of LV failure
P_{AF_T} = Probability of abort failure conditional on LV failure
P_{FP} = Probability of false positive LV failure indication
P_{AF_F} = Probability of abort failure conditional on a false positive LV failure indication

* LOC includes LOM in this analysis.
† However, for abort lead times that exceed the nominal 3.0 second burn time of the abort motor, it is possible that a still-

The probability of LV failure, P_{LVF}, was chosen equal to 0.003, which is considered to be in the range of ascent failure probabilities that have been calculated in recent years for human-rated LVs, such as those calculated in[3].

To calculate abort failure probabilities P_{AF_T} and P_{AF_F}, the Dynamic Abort Risk Evaluator (DARE)[4] was used. DARE models abort failure probability conditional on the occurrence of any of a number of user-selectable LV failure modes occurring at a user-specified time into ascent. For this analysis, DARE was modified to generically model abort failure probability given generic LV failure (which includes the potential for catastrophic LV failure and the resulting blast overpressure and fragmentation). In the case of a false positive, DARE was run using a failure mode with a benign abort environment (i.e., no overpressure or LV fragmentation).

In order to provide a quantitative basis for decision-making, utilities were defined for each of the three possible outcomes: a utility of one was assigned to OK; a utility of zero was assigned to LOC; and a utility of 0.9 was assigned to LOM. The basis for the OK and LOC utilities is convention – the range of utilities spans the interval [0, 1]. The LOM utility was chosen as 0.9 to reflect the value that it is much worse to lose the crew (and mission) than it is to lose the mission alone. Classical decision theory holds that the best decision is that which maximizes expected utility, E[U]. In the event tree of Figure 1,

$$E[U] = 1*P(OK) + 0.9*P(LOM) + 0*P(LOC) = P(OK) + 0.9*P(LOM) \quad (4)$$

III. THE DECISION

In DARE, one of the driving factors of abort effectiveness is abort lead time, i.e., the time interval between abort initiation and the production of any adverse abort environment by the failing launch vehicle. A longer lead time typically results in a greater distance between the crew vehicle and the point of origin of any blast wave or fragmentation cloud, thereby reducing any overpressure stress on the crew vehicle, as well as the probability that the crew vehicle will be hit by a launch vehicle fragment, leading to increased overall abort effectiveness.[†]

The decision used for this study is that of integrated vehicle health monitoring (IVHM) system selection, where the choice is between two options:

- Option A, which has a relatively low false positive error rate (P_{FP} = 0.001) but is believed to provide relatively short abort lead times, leading to lower abort effectiveness given launch vehicle failure; and

- Option B, which has a relatively high false positive error rate (P_{FP} = 0.01) but is believed to provide longer abort lead times and higher abort effectiveness.

As configured for this study, DARE models abort lead time as a random variable that varies uniformly over the interval $[\mu_{ALT} - (\ln(\mu_{ALT} + 1))/2, \mu_{ALT} + (\ln(\mu_{ALT} + 1))/2]$, where the mean abort lead time μ_{ALT} is epistemically uncertain over the interval $[\mu_{MIN}, \mu_{MAX}]$. In other words, abort lead time uncertainty has both an aleatory and epistemic component, with abort lead time varying uniformly from trial to trial around some unknown mean value, μ_{ALT}, that is believed to lie somewhere between zero and a user-specified upper bound, μ_{MAX}. Figure 2 illustrates the situation. Option A was assumed to have a μ_{ALT} range of [0s, 3s], whereas Option B was assumed to have a μ_{ALT} range of [1s, 4s].

IV. THE ANALYSIS

Figure 3 shows graphs of E[U] over the ranges of possible values of μ_{ALT} for Options A and B. DARE nominally treats μ_{ALT} as being uniformly distributed between μ_{MIN} and μ_{MAX}. To investigate the implications of introducing an interval characterization of uncertainty into an otherwise distribution-based model, the distribution for μ_{ALT} was substituted by an interval, with no commitments concerning where μ_{ALT} is likely to lie within the interval. Calculations of E[U] were made for both the wholly distribution-based nominal ascent model, as well as for the modified model where μ_{ALT} is treated as an interval, and the role of the results to support decision-making is discussed.

[†] However, for abort lead times that exceed the nominal 3.0 second burn time of the abort motor, it is possible that a still-accelerating LV can close the gap between the LV and the aborting crew vehicle, increasing the probability of abort failure.

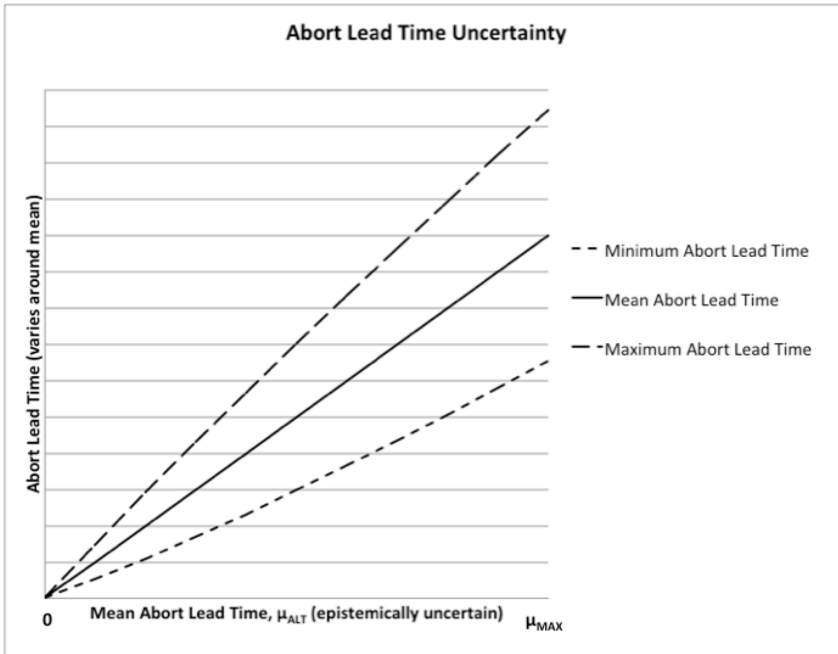

Fig. 2. Abort Lead Time Determination

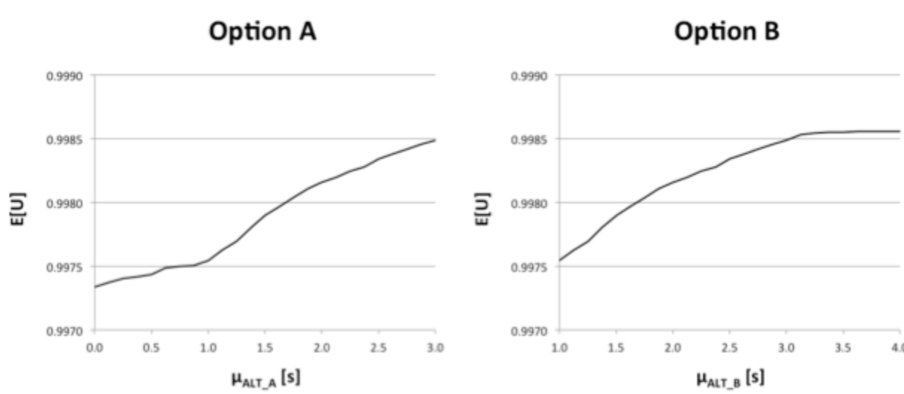

Fig. 3. E[U] vs. μ_{ALT} for Options A and B

IV.A. Modeling μ_{ALT} as a Uniform Distribution

The case where μ_{ALT} is modeled as a uniform distribution is consistent with a conventional treatment of epistemic uncertainty in probabilistic analysis. Distributions for P_{AF} can be calculated that incorporate all the epistemic uncertainties in the abort model, including abort lead time uncertainty, as shown in Figure 4 for P_{AF_T}.

More importantly, the full ascent model can be used to calculate point values for E[U] for each of the two options. As mentioned previously, classical decision theory holds that the best decision is that which maximizes E[U]. The results are shown in Figure 5. The points labeled "Mean" are the values for E[U] for Options A and B, respectively. The fact that E[U] is greater for Option B than for Option A means that Option B is the better decision, given the current state of knowledge.[‡,§]

[‡] This study does not make any distinction between "risk-based" decision-making, in which the results of a risk analysis are used directly to make a decision, versus "risk-informed" decision-making where risk analysis are but one input to a subjective, deliberative decision-making process.

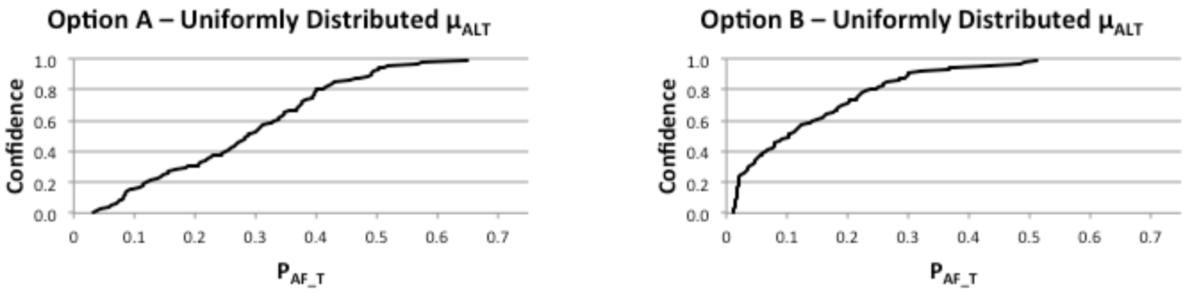

Fig. 4. P_{AF_T} Distributions for Uniformly Distributed μ_{ALT}

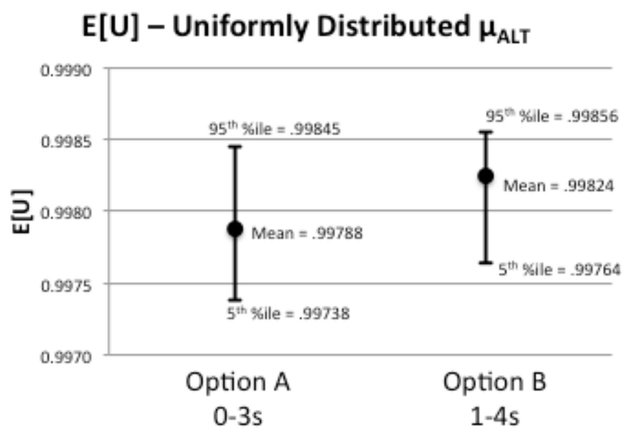

Fig. 5. E[U] for Options A and B in the Case of Uniformly Distributed μ_{ALT}

The error bars in Figure 5 show the variation of E[U] over the distribution for μ_{ALT}. In other words, for each option, the bars show, with 90% confidence, the range within which E[U] would be calculated if μ_{ALT} were known with certainty. The considerable overlap of these ranges between the two options suggests that although selecting Option B is a rational choice given current knowledge, it is not a *robust* choice, and it might be advantageous to obtain more information about μ_{ALT} in order to reduce the overlap and the associated probability that Option A is actually the better choice.

Accordingly, some investigation is made into the value of obtaining additional information about mean abort lead time prior to selection of IVHM system. Specifically, this analysis investigates the value of perfect information (VOPI), i.e., the value, in terms of expected utility, of knowing with certainty what the expected lead times are for the two options.[**] The situation is shown graphically in Figure 6. In the figure, the curves of Figure 3 have been projected over the μ_{ALT_A} - μ_{ALT_B} plane to create two intersecting surfaces. It can be seen that there are regions over the μ_{ALT_A} - μ_{ALT_B} where which Option B has the higher expected utility, and regions where Option A has the higher expected utility. If the joint distribution $f(\mu_{ALT\_A}, \mu_{ALT\_B})$ has mass over both these regions then the decision is not completely robust. Specifically, the decision to select Option B might yield an inferior system, if the true values of μ_{ALT_A} and μ_{ALT_B} are on the right-hand side of the domain, e.g., at 2.5s and 0.5s, respectively.

[§] The relatively small difference (~2%) between the minimum and maximum utilities on this graph is somewhat misleading. Because the consequences of failure (i.e., LOC, LOM) are so dire, there is a high desire to avoid these outcomes and even small probabilities are keenly felt. As such, the situation might be best appreciated in terms of *disutility*, i.e., 1 minus utility, which in this case varies by a factor of three along the vertical axis.

[**] This study does not concern itself with the value of imperfect information, nor with the cost of obtaining additional information, perfect or otherwise.

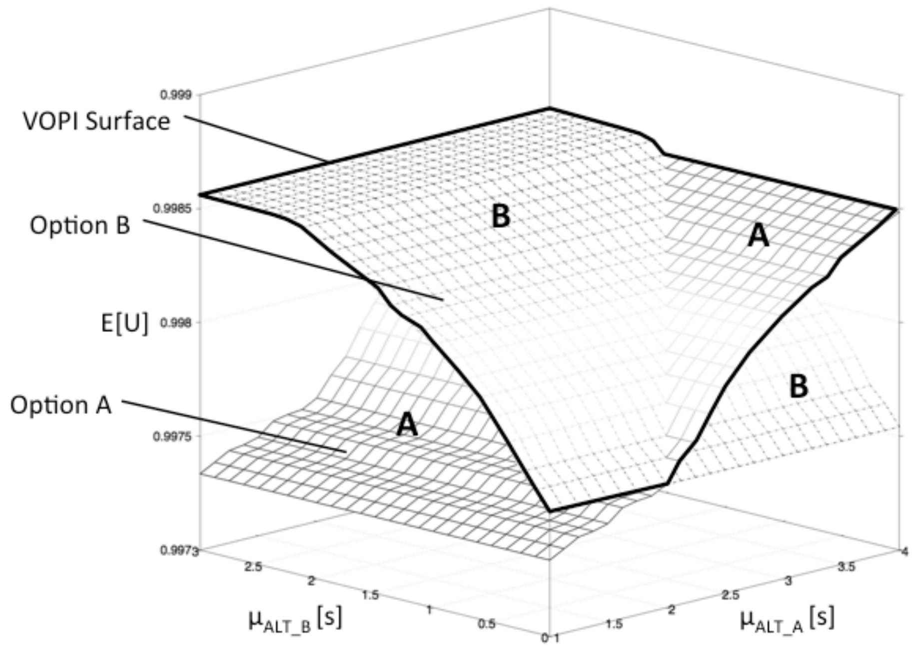

Fig. 6. The "VOPI Surface"

Given perfect information about μ_{ALT_A} and μ_{ALT_B}, it would be possible to use Figure 6 to pick the option with the highest value of E[U] above the point (μ_{ALT_A} ,μ_{ALT_B}). We define the "VOPI surface" to consist of all such values of E[U], i.e., it is the outlined upper surface of the figure, which is a composite of the A and B surfaces. The value of perfect information about μ_{ALT_A} and μ_{ALT_B} is therefore the difference between the expected utility of the VOPI surface and the expected utility of the Option B surface:

$$\begin{aligned} \text{VOPI} &= E[U]_{\text{VOPI Surface}} - E[U]_{\text{Option B Surface}} \\ &= \iint d\mu_{ALT\_A}\, d\mu_{ALT\_B}\, \{E[U]_{\text{VOPI Surface}}(\mu_{ALT\_A}, \mu_{ALT\_B}) \\ &\quad - E[U]_{\text{Option B Surface}}(\mu_{ALT\_A}, \mu_{ALT\_B})\}\, f(\mu_{ALT\_A}, \mu_{ALT\_B}) \end{aligned} \tag{5}$$

If this value exceeds the cost of obtaining the information, then the rational decision is to defer IVHM selection until after the information has been obtained, at which point a perfectly robust decision can be taken.

IV.B. Modeling μ_{ALT} as an Interval

Under an interval characterization of μ_{ALT}, the preceding analysis can be mirrored but the implications for decision-making are altered. Without probability distributions for μ_{ALT_A} and μ_{ALT_B} (or, more accurately, without a joint distribution $f(\mu_{ALT\_A}, \mu_{ALT\_B})$), E[U] cannot be calculated and there is not necessarily a rational basis for decision-making.

Figure 7 shows abort failure probability results for Options A and B, analogous to Figure 4 above, but now expressed as probability boxes (p-boxes) rather than probability distributions. These p-boxes admit the possibility of any monotonically increasing probability distribution spanning a confidence from zero to one that fits within the box, which is equivalent to admitting the possibility of any distribution for μ_{ALT} that fits within the interval [0, μ_{MAX}]. Similarly, an expected utility treatment of the options is possible, but each option can only be characterized as an interval over E[U], rather than a single value. Thus, the analogue of Figure 5 is Figure 8.

The most significant consequence of an interval characterization of E[U] is that it is insufficient for optimal decision, i.e., neither option maximizes E[U]. Instead, a different decision rule must be found. One such decision rule, which is consistent with a risk-averse attitude, is to select the option for which the minimum possible expected utility is highest. Such a *maximin* decision rule focuses on minimizing the potential downside of a decision, without regard for the magnitude of the potential upside. The maximin rule may be most appropriate when striving to achieve some threshold utility, such as when some threshold level of performance is required of a system, so long as the minimum performance exceeds the required level. Other decision rules are possible, and can lead to different decisions. A decision rule that reflects a risk-seeking attitude might be to select the option with the highest maximum possible expected utility. It just so happens that Option B would be selected under both of these rules, but if Option B had the smaller maximum E[U] then Option A would be selected under the aforementioned risk-seeking decision rule.

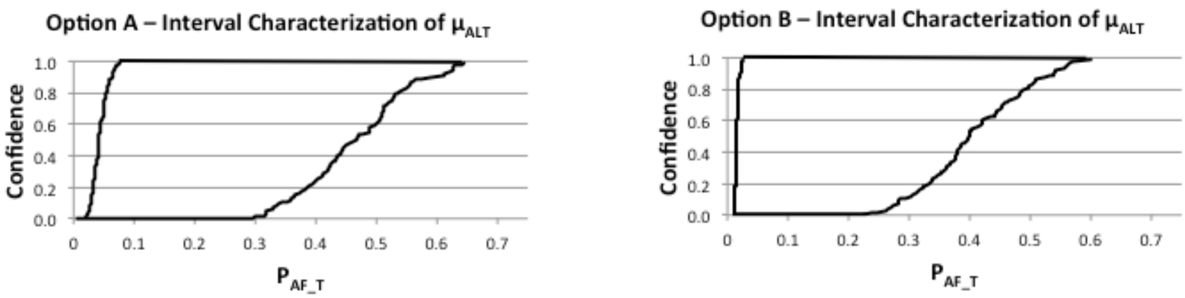

Fig. 7. P_{AF_T} P-Boxes for Interval-Characterized μ_{ALT}

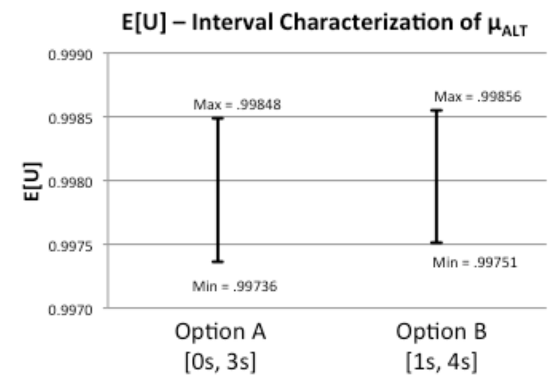

Fig. 8. E[U] for Options A and B in the Case of Interval-Characterized μ_{ALT}

As in the case of the wholly-distribution-based model, a robust decision cannot be made due to the overlap of the E[U] intervals, so it is worth examining the extent to which additional information might be of value. Not surprisingly, given the absence of a joint distribution over (μ_{ALT_A}, μ_{ALT_B}), it is not possible to calculate the value of information, perfect or otherwise. Instead, like E[U] generally, the value of perfect information can be characterized as an interval. Figure 9, which is the analogue of Figure 6, illustrates the situation.

First, let it be stipulated that under the current state of knowledge, Option B is the preferred option (e.g., based on a maximin decision rule). It is easy to see that the VOPI is minimized for any point (μ_{ALT_A}, μ_{ALT_B}) where Option B has a higher E[U] than Option A, such as on the left-hand side of the figure. Here, the VOPI is zero since additional information would not change the decision. The VOPI is maximized at the point where the difference between the E[U] of Option A and that of Option B is greatest, namely at (3s, 1s) at the far right corner of the figure. In this case, the VOPI is $E[U]_A - E[U]_B$. Thus, the VOPI is expressed as the interval $[0, \max[E[U]_A - E[U]_B]]$.

If, for any two options *i* and *j*, $E[U]_i$ and $E[U]_j$ are functions of the same interval-characterized uncertain parameter *q*, then the lower limit on the VPOI interval can be non-zero. This is seen in Figure 10, in which a maximin decision rule is applied to a choice between three options. Option C is the preferred option under the decision rule, but nowhere does it have the maximum E[U]. Consequently, the minimum VOPI is positive and the VOPI interval is expressed as:

$$\begin{aligned} VOPI &= [VOPI_{MIN}, VOPI_{MAX}] \\ &= [\min[\max[E[U]_X] - E[U]_P], \max[\max[E[U]_X] - E[U]_P]] \end{aligned} \tag{6}$$

where $\max[E[U]_X]$ is the maximum E[U] from among the non-preferred options at the point of comparison to $E[U]_P$, and $E[U]_P$ refers to the preferred option.

A situation where $VOPI_{MIN} > 0$ is one where the preferred option would definitely not be selected under perfect information, and can therefore be considered a kind of hedge.

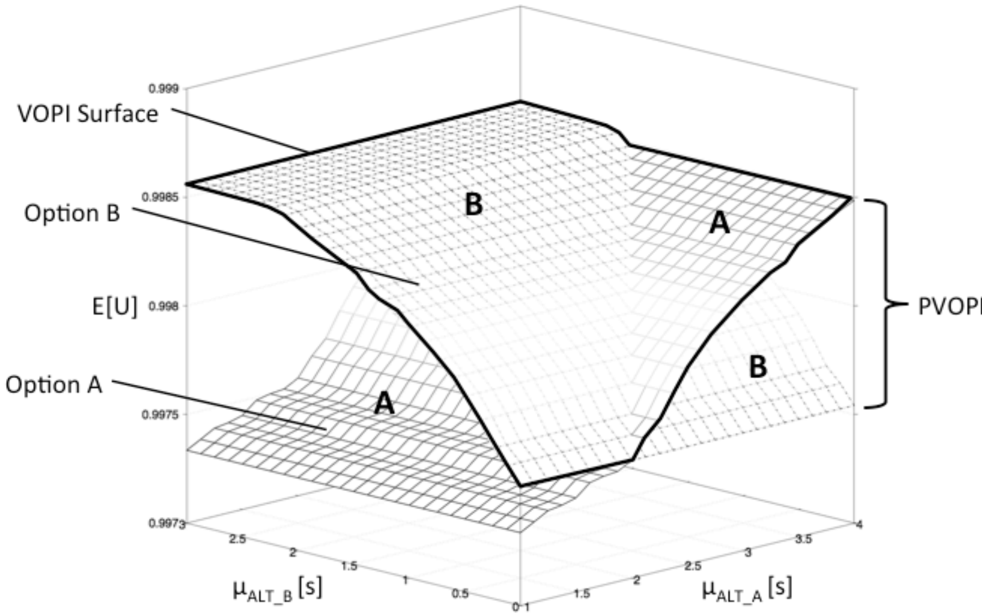

Fig. 9. The Potential Value of Perfect Information (PVOPI)

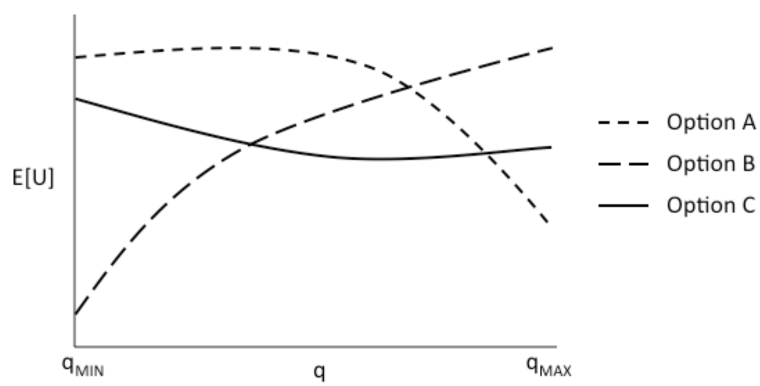

Fig. 10. A Decision Situation where $VOPI_{MIN} > 0$

In cases where $E[U]_i$ and $E[U]_j$ are subject to separate, uncorrelated uncertainties, then the situation is more like that of Figure 9 and it would seem that reasonable preferred option must have the greatest E[U] *somewhere* over the (joint) interval, in which case the minimum possible VOPI is zero. Then the maximum possible VOPI is all that is needed to specify the interval, and we can define the *potential* value of perfect information (PVOPI) as that maximum:

$$PVOPI = \max[E[U]_A - E[U]_B]. \tag{7}$$

ACKNOWLEDGMENTS

The authors would like to thank the NASA Office of Safety and Mission Assurance (OSMA) for their sponsorship of this study.

REFERENCES

1. GUARRO, "Open Conceptual Questions in the Application of Uncertainty Analysis in PRA Logic Model Quantification," ASCA Inc., Redondo Beach, CA, 2014.
2. FERSON et al, "Constructing Probability Boxes and Dempster-Shafer Structures," Sandia National Laboratories, Albuquerque, NM, 2003.
3. "NASA's Exploration Systems Architecture Study," NASA, Washington, DC, 2005.
4. EVERETT, HALL, "Dynamic Abort Risk Evaluator – Analysis Process Document," SAIC, New York, NY, 2007.