An Example of Decision Management

David G. Ullman, Robust Decisions Inc.

Abstract

Many decisions are based on simulation and test results. In spite of these modeling efforts, these decisions sometimes lead to poor products or business decisions. As an example, in this white paper, a decision is made about a concept for an Unmanned Combat Air Vehicle (UCAV) early in its design process. Simulation results initially show that a decision should be made to refine the concept. However, simple additions to the simulations begin to reveal potential problems with this configuration.

This problem is used to demonstrate the features of decision management. Decision management techniques will allow the uncertainty and inconsistency in the information to be taken into account and expose weaknesses in the proposed concept. Further these techniques will help integrate team evaluation and the consistent comparison between competitive concepts.

Decision management techniques used in this example can be applied to any technical, business, portfolio management, scenario planning or human relations issue.

Introduction

Decisions are a daily part of product and business development. Inherent in each decision are risks – risks that the choice made is not the best based on the information available, and risks that once a decision is made, the outcome will not turn out as expected.

The goal of decision management is to use the information available to make the best possible choices with known, expected risks, within time and resources available. This paper will show by example a recently developed decision management methodology that helps measure decision risks as part of the decision making process.

Although the example problem is for a product development decision, the methodology can be applied to IT portfolio management, outsourcing decisions (e.g. vendor selection and partner selection), personnel actions, business strategies, or any other situation with the need to evaluate and compare alternative outcomes. For each, the goal is to make the best possible decision with known levels of risk.

The term "risk" is usually not well defined. In this paper we have a specific, calculatable method for determining what is commonly referred to as risk. Actually, we divide "risk" into two measures. Referring back to the first paragraph: (1) "the choice made is not the best based on the information available" and (2) "once a decision is made, the outcome will not turn out as expected".

The first measure of risk is the level of satisfaction with the choice made. This is both a relative and absolute measure. Relatively, how much is the satisfaction with the option chosen better than other options considered, and absolutely, is the satisfaction high enough to warrant commitment to it. If the satisfaction with an option is low, the risks in selecting it are high. If the satisfaction with an option is not much better than other options then the risks are high that the "best" option has not been chosen. Throughout the paper, this measure of risk is referred to as "satisfaction."

The second measure of risk relates how much the satisfaction might degrade if the option does not perform as anticipated. In other words, how much worse might things be than initially thought? Throughout this paper we refer to this measure as "decision risk" or just "risk."

These definitions lead to two questions that recur throughout decision-making:

1. Which option is best?
2. What are the risks inherent in it?

These two questions will be answered in the example problem through a series of Evaluations. Each Evaluation adds reality to the situation, increasing the understanding of the satisfaction and risk, and increasing the potential of making a good decision.

The evaluations help answer the following questions.

Evaluation 1: How does **uncertainty in goals and evaluation information** affect the satisfaction and risk?

Evaluation 2: How do **different views about what is important** affect risk?

Evaluation 3: How can **different options be compared** using measured satisfaction and risk?

Evaluation 4: How can the **opinion and knowledge of the team members** be pooled even though they are distributed in time and location, represent different organizations, and even though they use different evaluation measures?

Evaluation 5: How can we insure **uniform treatment of technical, business and other important measures** when making a choice?

Evaluation 6: **What should be done next** in the decision making process to insure that the best choice is being made?

Evaluation 7: How can **satisfaction and risk be tracked as the project evolves**?

And, how can we do all this within time and resource limitations. The software *Accord*™[1] will be used to support these Evaluations.

[1] *Accord* can be downloaded from www.robustdecisions.com. The Demo version of the software contains the data file for this example problem.

Example Problem

This problem is based on one used by Dr. Tyson Browning of Lockheed Martin Aeronautics Company[2] for the conceptual design of an Unmanned Combat Air Vehicle (UCAV). **The goal is: choose a UCAV concept early in the design process for refinement that has a good chance of successful development.**

The performance targets used for the UCAV are given in Table 1. All the targets are of the form more-is-better. In other words, a UCAV's ability to carry an 800 lb payload is the target. Anything less than that decreases satisfaction with the concept and anything more than that may not be necessary.

Measure	Target	Units	Comments
Payload	800	Lbs	
Range	725	Nms	
Reliability	5000	Hrs	Mean time between failures (mtbf)
Detectability	1.5	-	Ratio with previous vehicle

Table 1. Initial targets

Before exploring this problem through the seven Evaluations, consider a baseline for a single concept, Concept 1. Initial evaluation of this concept has predicted the performance shown in Table 2.

Measure	Target	Units	Simulated Results
Payload	800	Lbs	850
Range	725	Nms	750
Reliability	5000	Hrs	5100
Detectability	1.5	-	1.55

Table 2. Initial performance estimates for Concept 1

Since all the predictions show that the targets are met, then this is an acceptable concept – we are 100% satisfied with it. However, this conclusion assumes that there are no errors or other uncertainties in the simulations and that the four measures fully characterize the vehicle.

This type of deterministic analysis is often the only form of evaluation that is undertaken. In order to include the effect of uncertainty, sometimes a sensitivity analysis is undertaken after-the-fact to explore the effect of some parameters on the results. As shown in the next Evaluation, this analysis comes too late.

---

[2] Browning, Tyson R., John J. Deyst, Steven D. Eppinger, and Daniel E. Whitney (2002) "Adding Value in Product Development by Creating Information and Reducing Risk," IEEE Transactions on Engineering Management, Vol 49, No 4, Nov 2002. pp443

Further, this deterministic analysis has not given sufficient information to answer any of the questions asked in the introduction. A measure titled, “risk”, could be added to the list of four measures with the goal of making risk as small as possible. However, with this addition it is not clear whether this risk measures the possibility of not making the range, payload or some other combination of measures.

Evaluation 1: How does <u>uncertainty in the goals and evaluation information</u> affect the satisfaction and risk?

More realistically, there are many uncertainties in every simulation, especially of early concepts. A more robust approach is to recognize these uncertainties and integrate them into the analytical effort. This approach will also allow us to better estimate the satisfaction and monitor the risks associated with the concept. In this Evaluation we will add uncertainties in the analysis and targets to the baseline. These will allow us to better calculate the satisfaction and risk associated with this concept at this point in its development.

We represent uncertainty by collecting high, low, and most likely values for each measure. These values can result from simulations, experiments or best-guess estimates. Assume that the most likely values for the measures are the evaluation results given in Table 2, and the high and low represent 3 standard deviations ($\pm 3\sigma$) from this value. Note that the distribution may not be symmetrical about the most likely value. For Concept 1, the values are shown in Table 3:

Measure	Target	Units	Low	Most Likely	High
Payload	800	Lbs	750	850	1000
Range	750	Nms	650	750	900
Reliability	5000	Hrs	4500	5100	6000
Detectability	1.5	-	1.2	1.55	2.0

Table 3. Performance estimates for Concept 1

For the Reliability criterion, the evaluation distribution based on these values can be estimated to look as shown in Figure 1.

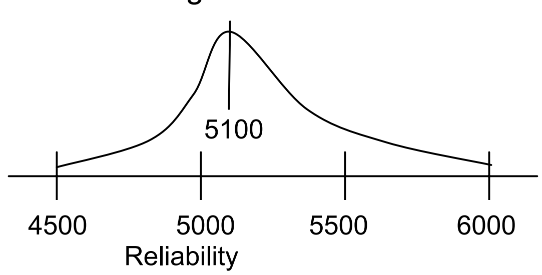

Figure 1. Distribution for Reliability

In this figure about 40% of the area under this curve is below 5000 hrs – not meeting the target. Thus, at this point in the evaluation of Concept 1, even though the most likely estimate is satisfactory, 5100 > 5000 hrs, there is about a 40% chance that the target will not be achieved. This distribution is important in the calculation of satisfaction and risk.

The target for the reliability was given as 5000 hrs. A better representation for a target is to represent it as a utility curve. A simple utility curve for reliability is shown in Figure 2. The curve shows that the ideal reliability is 5000 hrs (any reliability above this is ideal) and the minimum target value is 4500 hrs (any reliability below this is unacceptable). The two points are connected with a straight line. More complex fits are possible, but any more sophistication is not warranted.

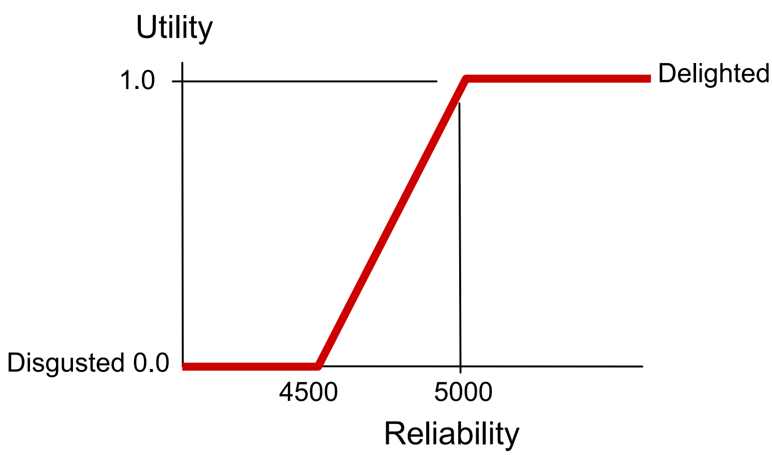

Figure 2. The utility curve for reliability

The utility values given on the vertical axis translate the reliability into a number between 0 and 1 that reflects the worth of the reliability value. A reliability of 4500 hours is worth nothing and a value above 5000 hours has full utility of1.0.

If both the ideal target and the minimum acceptable are the same, i.e. 5000 hrs, then this model has reverted to the simple, single valued target as in Table 1– anything below 5000 is not acceptable and anything above is.

In the example problem, the ideal target and the minimum acceptable values for each of the measures is given in Table 4. A utility curve similar to the one for reliability, given above, can be drawn for each measure.

Measure	Ideal Target	Minimum Target	Units
Payload	800	600	Lbs
Range	750	700	Nms
Reliability	5000	4500	Hrs
Detectability	1.5	.9	-

Table 4. Two valued targets

The relationship between the evaluation and the targets can be better visualized on number lines. Those in Figure 3 are screen captures from *Accord*™. The thumbs up and down show the ideal and minimum acceptable targets respectively. The large dots show the most likely estimates and the brackets show the high and low estimates.

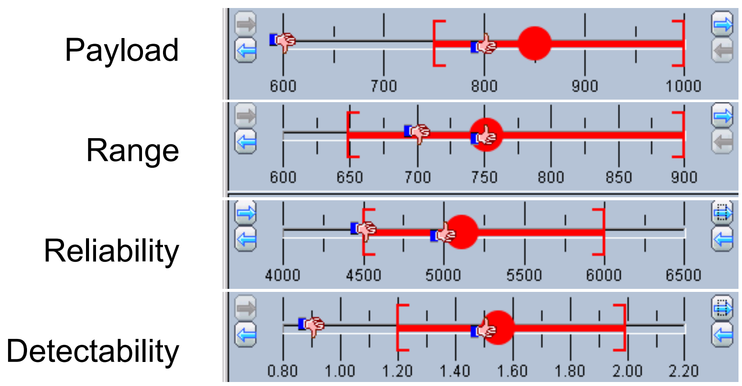

Figure 3. Evaluation of Concept 1

With high uncertainty in the evaluations, it is obvious that Concept 1 may not be as good as first thought.

In *Accord*™ the measure of how well a concept meets the criteria is referred to as the “satisfaction” – where 100% satisfaction means that the concept is predicted to fully meet the criteria. The lower the satisfaction, the less likely the concept is to meet the criteria. Technically, satisfaction is the common name for expected utility. Using *Accord,* this concept has a satisfaction value of 86%. This result assumes that you can trade off good performance in one area against poor performance in another (e.g. a high payload can compensate for a poor range). This may not be the best assumption and there is more discussion on different tradeoff schemes in Evaluation 3.

Additionally, the risk in selecting this configuration can be found. In *Accord* **risk is the expectation of the configuration not meeting the criterion as well as anticipated.** Although we expect a satisfaction of 86%, *Accord* calculates that, due to the uncertainty, this value may be as low as 72% - there is a risk of 14% (86% - 72%) that things will not turn out as expected. Risk is found by combining the probability of being below the expected satisfaction, multiplied by the consequence of missing the ideal target. This product is used to compute how much lower the expected value might become if things do not go as expected.

Further, the risk is computed for each criterion. For this example, the composition of the 14% risk is shown in Figure 4. There is the risk that range alone can lower the satisfaction by 8% (57% of 14%).

Since risk is computed for each measure, a consistent understanding of how the potential of not meeting each ideal target can be easily developed. This will be reinforced in Evaluation 5 where some of the measures will be a mix of quantitative and qualitative.

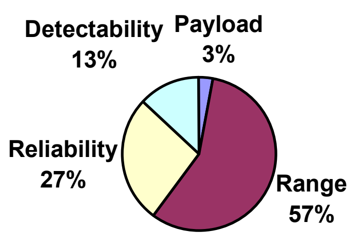

Figure 4. Risk contribution

In conclusion, even though baseline evaluation showed Concept 1 was satisfactory, it is only 86% satisfactory and the risk in selecting it is substantial at this time. Further, the evaluation shows that 57% of the risk is with the ability to meet the range requirement. These results assume that satisfaction with one criterion can compensate for poor performance in another. This is a weak assumption at best and it will be explored in Evaluation 2.

Evaluation 2: How do <u>different views about what is important</u> affect risk?

The first evaluation assumed that a single person provided information for judging the concept. In most situations there is a team involved working to evaluate the options and which is accountable for the decision. The involvement of a team introduces three new dimensions to decision management:

- Differing visions about what is important

- A range of evaluation results
- A richer set of criteria.

The first of these is addressed in this section with a range of evaluation results discussed in Evaluation 4 and a richer set of criterion added in Evaluation 5.

In the previous Evaluation, the Range and Reliability criteria had the highest risks associated with them. One question is, how important is Range relative to Reliability, i.e. must both targets be met to have a satisfactory concept? The answer to this question may vary with whom you ask and their functional responsibility within the organization. There are really two questions posed, how relatively important are the criteria to each other, and which criteria absolutely must be met. Each will be addressed by considering various viewpoints about what is important (Table 5).

Viewpoint	Assumptions	Payload	Range	Reliability	Dectability	Satisfaction	Risk
1	All can be traded off	.25	.25	.25	.25	86%	14%
2	All can be traded off but payload & range are more important	.35	.15	.20	.30	89%	11%
3	Range target must be met	.35	must	.20	.30	64%	38%
4	Detectability target must be met	.35	.15	.20	must	82%	18%
5	Range and reliability are must	.35	must	must	.30	57%	45%
6	All are must	must	must	must	must	54%	48%

Table 5. Various Importance Viewpoints

In Table 5 there are six viewpoints taken:

Viewpoint 1: Poor performance on one criterion can be traded off against good performance on another, and all of the criteria are of equal importance. These assumptions were used in Evaluation 1.

Viewpoint 2: Payload and Range are more important than Reliability and Detectability. From this viewpoint, Concept 1 is 89% satisfactory and the risk is lower than for the first option.

Viewpoint 3: The performance on all the criteria can be traded off except for Range, which must be met. From this viewpoint the satisfaction with Concept 1 is much lower (64%) and the risk higher as little is known about the range (See Figure 3).

Viewpoint 4: Detectability target must be met. The effect on the satisfaction and risk is smaller than with Viewpoint 3, as detectability is much better known (See Figure 3).

Viewpoint 5: Reliability and Range must be met and so, Concept 1 looks weak with risk only 12% lower than the satisfaction.

Viewpoint 6: To be fully satisfied with Concept 1, viewpoint 6 requires that all the criteria must be met (i.e., The UCAV must carry at least 800 pounds <u>and</u> go at least 750Nm <u>and</u> have at least 5000 hours between failures <u>and</u> have a detectability ratio of at least 1.5). From this viewpoint Concept 1 has only 54% satisfaction. It is easy to see why the value is so low by looking at Figure 3 and realizing that there is significant probability that the concept will not meet the Range or Reliability targets. Further, Concept 1 is very risky. With a 48% risk, if everything goes wrong that can go wrong, then the satisfaction is only 8% (54%-48%). Unfortunately this viewpoint of the problem is probably the most accurate.

It is not clear which of these viewpoints to assume, or whether to assume some combination of them. There is no right answer! However, a methodology for managing these differences will be developed in Evaluation 5. Before addressing that issue, consider comparing other alternatives to Concept 1.

Evaluation 3: How can <u>different options be compared</u> using measured satisfaction and risk?

The goal of this exercise is not to just evaluate Concept 1; the goal is to choose the best from a set of proposed concepts (the goal of all decision-making). We are addressing the first question posed: Which alternative is best?

As an example, there are three options being evaluated: Concept 1 as described above, and two additional alternatives, Concepts 2 and 3. Evaluation of the Concepts is as in Figure 5. The bars show the low, most likely, and high values based on simulations and the solid and dashed lines show the ideal target and minimum target values.

Concept 2 is very similar to Concept 1 but the payload, range and reliability are better known and the detectability is not as high. Configuration 3 has better detectability and reliability and is weaker in range and payload. Which Concept should be chosen and which measures are worth spending additional time refining?

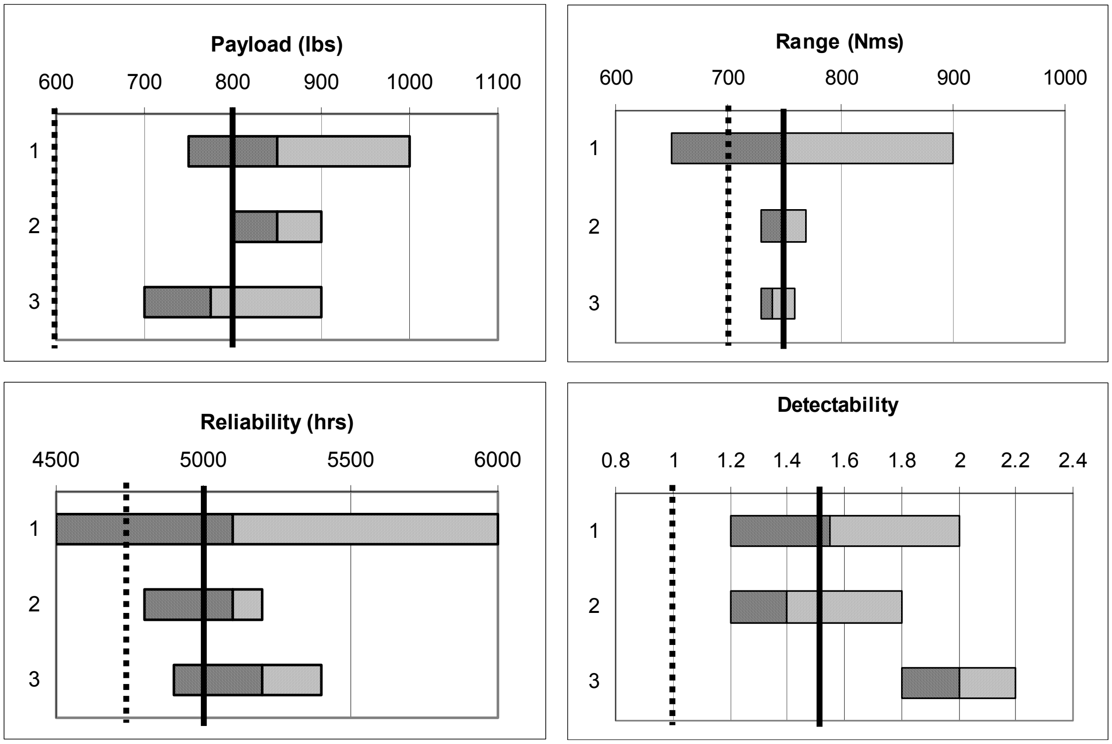

Figure 5. Performance estimates for three concepts: 1, 2, and 3
(━━ = ideal target , •••••• = minimum target)

If all the criteria are “musts” (all must be met), *Accord* calculates the satisfactions and risks shown in Figure 6. Here the first Concept has a satisfaction of 54% with a risk of 48%. These are the same values as on line six of Table 5. They imply that the satisfaction may be as low as 6%. Concepts 2 and 3, on the other hand, have much higher satisfaction and lower risk. With Concept 2, the only performance measure that improved Concept 1 was Detectability. A majority of the increase in satisfaction and decrease in risk came from improved performance estimates.

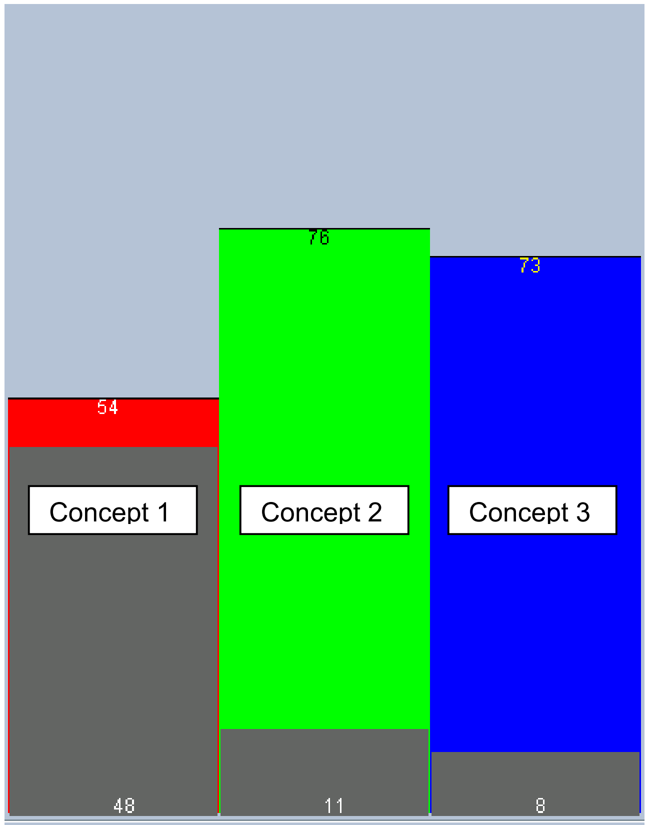

Figure 6. Satisfaction and Risk for three Configurations

Each of the risks is detailed in Figure 7.

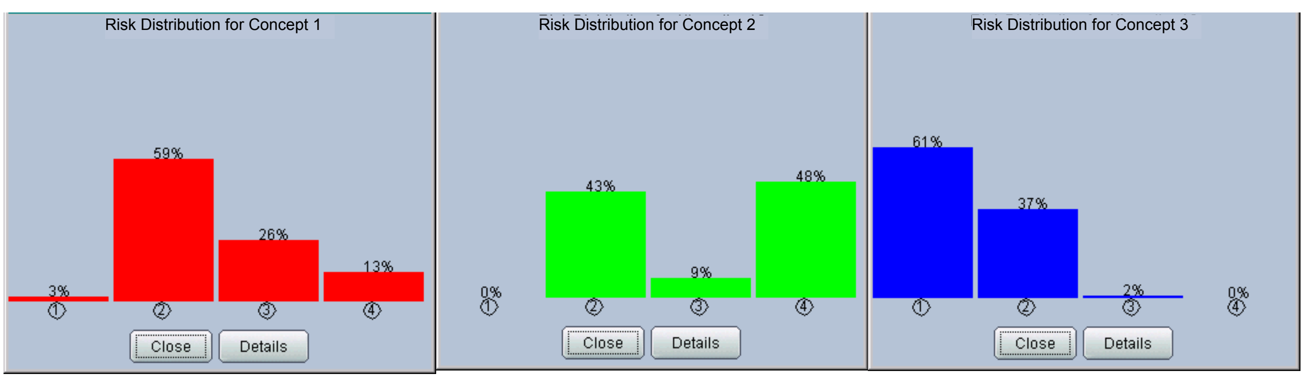

Figure 7. Risk analysis for the three concepts
1 = Payload, 2 = Range, 3 = Reliability, 4 = Detectability

It is clear that Concept 1 has the highest risk with range. For Concept 2 it is with both range and detectability. And, for Concept 3 it is the payload. These results may be discernable from the original data, but this is unlikely. The analysis performed here shows the results in a uniform and comparable manner that is easy to comprehend. Also, in Evaluation 6, this and other information can help determine what to do next: is it time to select one Concept and drop the others, or is more work needed to insure a good decision is being made.

Evaluation 4: How can the opinion and knowledge of the team members be pooled even though they are distributed in time and location, represent different organizations, and even though they use different evaluation measures?

There are two aspects to this Evaluation: differences in team member opinion and knowledge, and the fact that the team may be distributed in time and location. We will address them one-at-a-time.

As stated at the beginning of Evaluation 2, the involvement of a team introduces three new dimensions to evaluations:

- Differing visions about what is important
- A range of evaluation results
- A richer set of criteria

In Evaluation 2, we discussed multiple viewpoints about what is important, yet we did not resolve how to manage them. In this section we will address this issue and also a methodology for managing conflicting evaluations.

There are three ways to manage differences in what is seen as important:
1. Ignore the differences
Have you ever been in a meeting that seemed endless? Everyone is talking about the same issue, and the discussion keeps going around and around. One possible cause of this is a failure to manage the fact that what one person thinks is important is different from what others think is important. Consequently, the meeting ends only when people are out of time or patience, or the most powerful person in the room makes the decision. In any case, the result will have poor buy-in by the others and the decision will likely not be very robust.

2. Negotiate the differences
Rather than ignore the differences, a more refined approach is to negotiate and try to arrive at some agreement. The goal of negotiation is to try and make winners of all the critical stakeholders. Finding the intersection of criteria that all consider important (i.e., a win-win situation) insures agreement. If there is no intersection, then the goal is to find a relaxed set of weightings that still address the issue and lead to a win-compromise situation.

It is worth noting that some decision-support consultants say that the best way to treat differences is to come to an agreement with a single, common viewpoint. They say there should be a single company viewpoint about what is important.

Yet, trying to average viewpoints loses the richness of team members' responsibilities and opinions.

3. Honor the differences
The third approach to working with differences of opinion is to recognize and articulate these differences as a richness factor and then discover alternatives that are satisfactory regardless of them. With this view, the differences are honored and recognized rather than ignored or averaged.

To show how *Accord* can help honor different viewpoints, consider a team of three people. Assume they weight the importance of the criteria as shown in Table 6. These are three of the examples used in Evaluation 2.

Team member	Assumptions	Payload	Range	Reliability	Detectability
DuLon Jenkins	Range target must be met	.35	must	.20	.30
Sally Field	Reliability and detectability are must	.35	must	must	.30
David Smithers	All are must	must	must	must	must

Table 6. Team Members' Importance Weightings

Using the evaluation information in Table 5, and the weights in Table 6, the resulting satisfactions are shown on one bar chart in Figure 8. Here the satisfaction results are shown superimposed in order to show the variation cause by viewpoint variation. For Concept 1, the satisfaction based on one team member's viewpoint is 64%, that of another team member is 54% and the third is in between these two values. Concept 2 seems to be the leader, but these results are still based on rather sparse evaluation. This observation leads to the second effect of a team on evaluation - a range of evaluation results.

Typically, many team members contribute to the evaluation of alternative concepts. Part of the richness in this contribution, is that they may not always agree on how well alternatives meet the targets. For example, one team member may have computed the payload as shown in Figure 3 and reprinted in Figure 9a. However, another may have estimated the payload as shown in 9b and yet another as in 9c. Their estimates may be interpretations of the same data or they may rely on different simulations or other information. It may not be clear which, if any of them are correct – they are all estimates that should not be ignored.

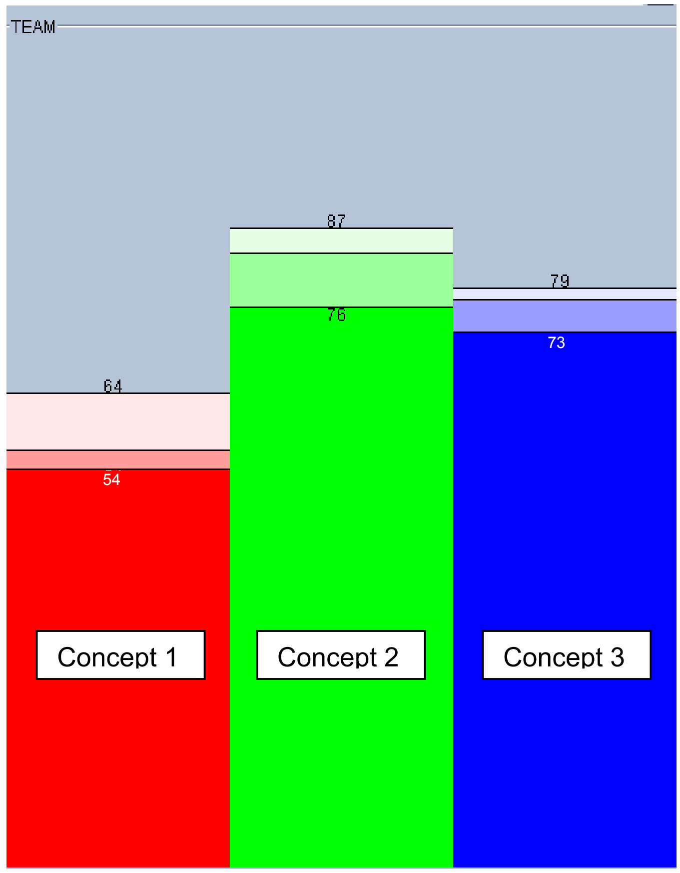

Figure 8: Team Satisfaction

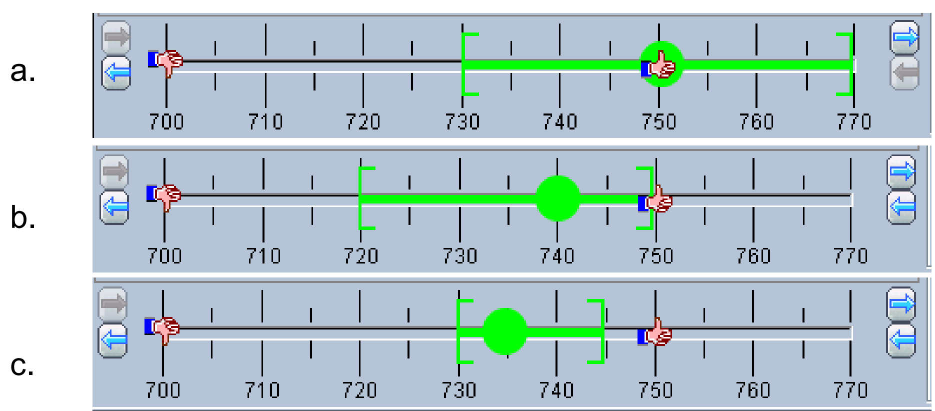

Figure 9. Evaluations for Range of Concept 2

Accord has the ability to combine these data into estimates of satisfaction and risk. For Concept 2, using all the evaluations in Figure 9, rather than only just the first one, lowers the range of satisfactions from 76-87 to 71-83. The methodology used to combine the evaluations is based on Bayesian methods and is detailed in another paper.[3] As will be seen in Evaluation 5, this capability is even more important when the criteria are not measurable.

In most critical decisions, the different visions about what is important and the range of evaluation results come from experts who are **distributed in time and location**. Flying everyone to a central location is costly and time consuming. Teleconferencing helps, and still information is often not structured to make this effective; and then there are time zone differences. *Accord* helps structure the information for efficient evaluation. It manages the collection and merging of information from either a centralized database or, if that is not available, by email attachments. In either case, once the alternatives and the criteria are defined, team members can input their weightings and evaluations whenever and wherever they are.

In the UCAV example, the team members were located at three different sites. As they evaluated the alternatives and saved their work, it was automatically merged in a database. The project manager then easily reviewed the shared team evaluation and directed the project based on the evaluation results.

Evaluation 5: How can we insure uniform treatment of technical, business other important measures when making a choice?

All of the measures used so far have been objective - quantitative. However, in most situations these are not the only measures used to make a decision. In making portfolio decisions - decisions on what products or projects to pursue - a majority of the key factors are not easily measured. For example, the book "Portfolio Management for New Products"[4] lists 19 measures grouped in 5 categories. The categories are:

- Probability of technical success
- Probability of commercial success
- Reward (to the company)
- Business strategy fit
- Strategic leverage

All key factors, with the exception of Reward, are subjective measures, based on estimates made by team members. These factors are based on the team's knowledge and educated guesses as to how well the proposed program will do.

For the UCAV example, additional important qualitative measures that might be important are:

- Probability of technical success (the different concepts use different technologies)

[3] http://www.robustdecisions.com/BayesianMethodDecisions.pdf
[4] Portfolio Management for New Products, R. G. Cooper, S.J. Edgett and E. Kleinschmidt, Addison Wesley, 1998.

- Ability of vendor to deliver
- Program complexity
- Consistency with current fleet
- Potential for extension and revision

Some additional quantitative measures that need to be considered are:

- Program cost
- Time to develop

Schedule and Cost measures can be evaluated as were the performance measures, but Ability of Vendor to Deliver and Probability of Technical Success are very hard to measure. It is easier to treat them qualitatively. At the same time these qualitative measures need to be combined with the quantitative measures in a uniform manner.

In *Accord,* evaluation of qualitative measures is accomplished using a Belief Map™. Belief Maps are described in detail at www.robustdecisions.com. A brief introduction is sufficient to explain how these new measures are evaluated.

Figure 10 shows the belief map for all three team members' evaluation of Concept 1's "ability of vendor to deliver." The resulting satisfaction and risk are based on the four performance measures; plus the two additional subjective measures.

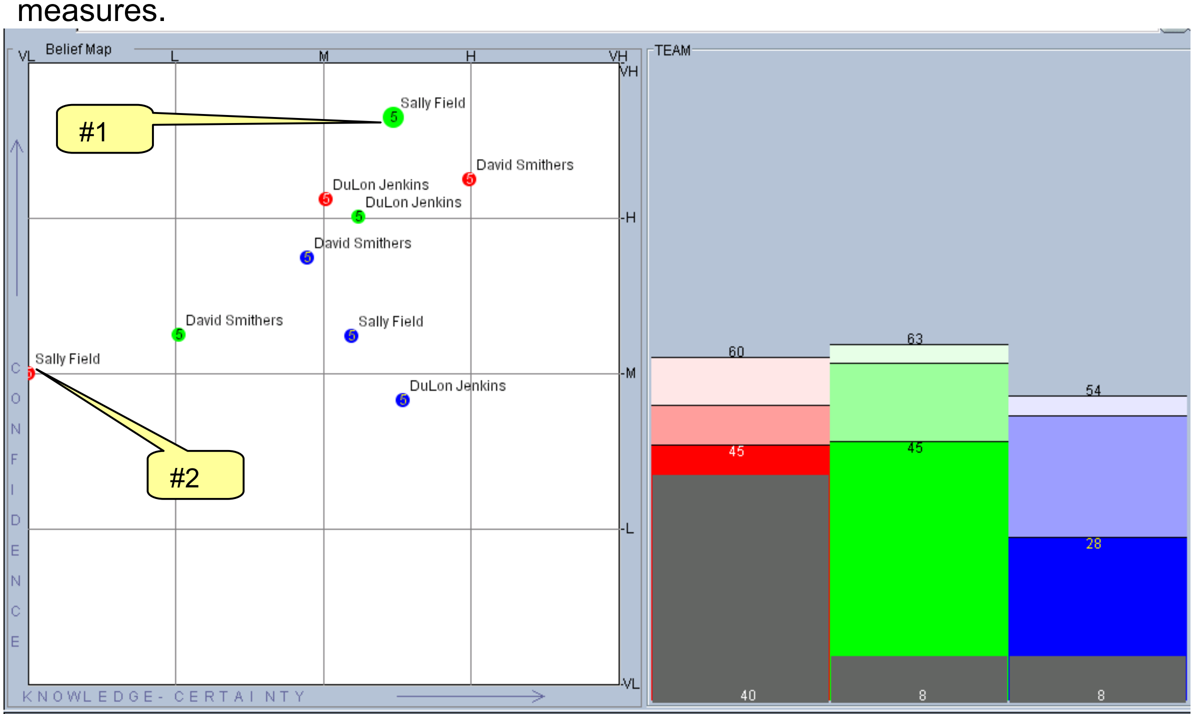

Figure 10. Belief Map and Satisfaction/Risk Results

Each member has evaluated on the belief map the ability of the three alternatives' vendors to deliver products. Each member places a point to signify his/her level of knowledge and confidence. For example, consider the point identified as #1. Sally Field has placed this green (for Concept 2) point to signify that she has high/very high confidence in Concept 2's vendor to deliver and that this is based on medium/high knowledge. *Accord* translates this location on the Belief Map into a probability that can be uniformly combined with those developed for the performance measures. With point #2, the default location for all points, Sally is saying that she has no knowledge on which to judge Concept 1. A point in this location has no effect on the results.

These additional evaluations change the satisfaction and risks for each of the alternatives. In Figure 10, Concept 2 has slightly better satisfaction than Concept 1, and much less risk. It has much better satisfaction than Concept 3. Is Concept 2 the one to choose? In the next Evaluation we will explore what to do next - the main point in studying the results.

Evaluation 6: <u>What should be done next</u> in the decision making process to insure that the best choice is being made?

One goal of decision management is to decide what to do next. In the words of Arthur C. Clark, "The only real problem in life is what to do next." Every decision (or procrastination in making one) costs time and other resources. Typically, what to do next is either: select an alternative and move to the next problem; generate more alternatives or refine those you have: refine the criteria used to evaluate the criteria; or discuss an evaluation and work toward consensus.

It is difficult, with so much information to see the benefit of doing more work on a concept versus the cost of doing so. However, *Accord* performs a value of information analysis that allows it to develop a "What to do next report". In Figure 11, this report details what further work is of the greatest benefit.

What to do next suggestions are ordered so that the first suggestion will have the highest benefit for the least effort. In the report, the evaluation that has the greatest potential for changing the results is to develop consensus on the Probability for Technical Success for Concept 1. This is followed by suggestions for further development of knowledge. The "what to do next" report is only blank when one alternative clearly dominates. Realistically however, one alternative usually dominates the others in its high satisfaction and low risk before this limit is reached.

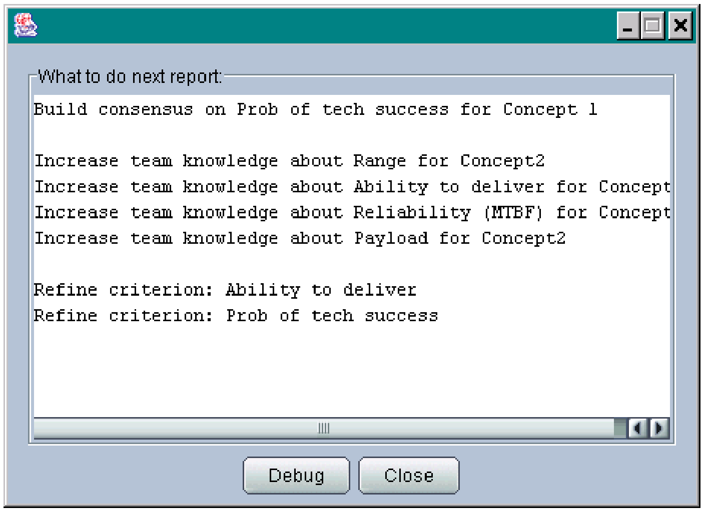

Figure 11. A what to do next report

Evaluation 7: How can satisfaction and risk be tracked as the project evolves?

Over time, more is learned about each concept. Figure 12 shows the performance evaluation information for Concept 1 at three different times in its development.

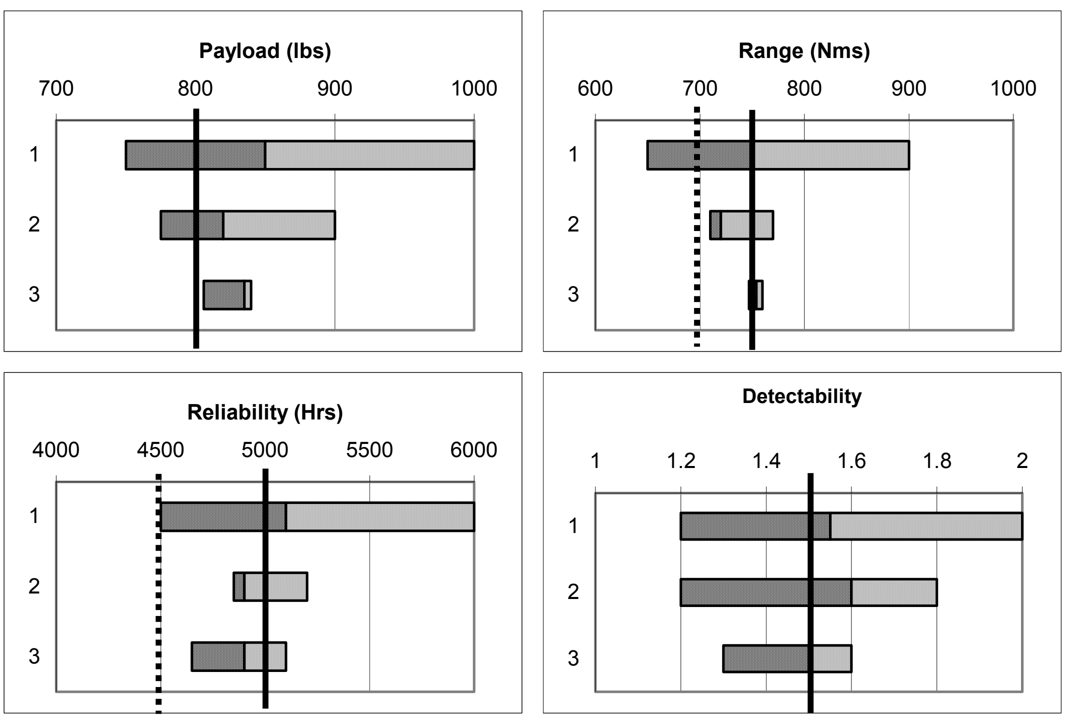

Figure 12. Evaluation change over time for Concept 1.

1= Concept design, 2 = Mid project, 3 = Release for production (solid line = ideal target, dashed line = minimum target value (off scale in some cases))

Some of the measures are brought to target with low uncertainty, for others the target is never met and they retain a significant level of uncertainty. The satisfaction and risk for Concept 1 over the period has followed the curves shown in Figure 13.

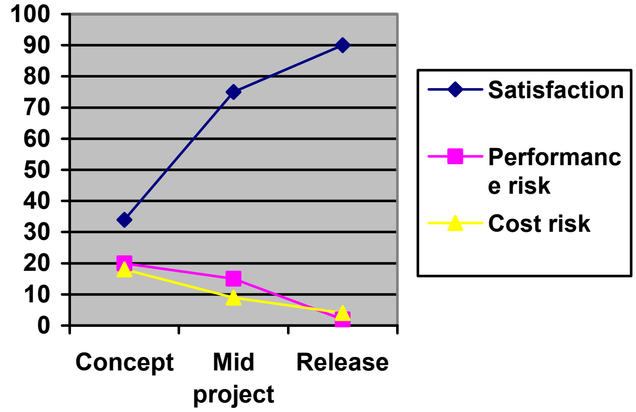

Figure 13. Change in Concept satisfaction and risk over time

Conclusion

At the beginning of this paper we stated two questions that recur throughout decision-making:

1. Which option is best?
2. What are the risks inherent in each option?

These two questions led to a series of Evaluations that led through uncertainty, risk, distributed team members with differences of opinion and knowledge, diverse criteria and project evolution, each of these adding reality to the decision made.

Even though the Evaluations shown in this paper may seem complex, when the decision management tools in *Accord are used*, they are made intuitive. Much of the work in this paper was defining criteria and targets, tasks that are necessary for any well managed project. Evaluation information from a team member only requires moving sliders on a number line or dots on a Belief Map. *Accord* manages all the details. In critical issues, whatever up-front work is necessary; Accord will generally save time and money later. In the time it took to read this paper, most of this analysis could have been accomplished using *Accord*.