

MULTIPLICATIVE TRADE STUDIES

James Eridon

General Dynamics Land Systems, Sterling Heights, MI

ABSTRACT

Combat vehicle design necessarily involves trade studies that attempt to balance varying performance criteria against associated burdens, such as weight, cost, and risk. Typically, trade studies quantify the value of different options by evaluating each on the individual criteria and then generating a weighted sum score. In the method described here the score is generated multiplicatively rather than additively. In addition, the importance of each evaluation criterion is used to generate the utility scores for that criterion. This improves the trade study process in two significant ways. First, making the overall score multiplicative greatly reduces the “compensation” problem, where good performance in some criteria can outweigh exceedingly poor performance in others. With a multiplicative method, only balanced solutions can score well. Second, using importance weights to establish utility scores for individual criteria simplifies the process, making it easier to conduct trades and evaluate the sensitivity of the results.

Citation: J. Eridon, “Multiplicative Trade Studies”, In *Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium (GVSETS)*, NDIA, Novi, MI, Aug. 10-12, 2021.

1. INTRODUCTION TO TRADE STUDIES

Trade studies are an important part of the design process, as decisions are made between competing priorities in an attempt to pick the “best” component or sub-system, all things considered. To assist in this process, formal methods have been developed which attempt to determine a composite score for different options, taking into account the key criteria, their relative importance, and the performance of different options against these metrics. However, it’s important to realize that the selection of an option with the best score is not, in fact, the goal of a formal trade study. Instead, the object is to identify the *reasons* for selecting one option over another. Any formal trade process should not claim to select the “best”

option for any given situation. Instead, it should help the designer determine why one selection is better than another. This goal has the advantage of not only showing which existing options may have good overall performance, but also identifying the sensitivity of the choice to the key criteria and their relative importance. Sensitivity analysis can help to identify potential design changes that could make any given option outstanding. Conversely, it also makes it possible to determine whether many potential options are so close in overall performance that the choice of any single one is not crucial.

Trade studies [1-5] are normally conducted in a formal manner in which a set of performance and burden criteria ($k=1\dots K$) are selected to evaluate a

range of potential options ($n=1\dots N$). The options might consist of different engines, for example, and the criteria might include horsepower, torque, maximum speed, weight, cost, and risk. A typical trade study proceeds by assigning an importance weight to each criterion (w_k) and generating a utility score ranging between 0 and 1 for each option on each criterion (U_{kn}) based on how well it performs relative to objectives. The overall score for a given option (S_n) is then calculated as a weighted average of the individual utility scores and the importance weights, so that we have:

$$S_n = \sum_{k=1}^K w_k U_{kn} \quad (1)$$

Importance weights are generally normalized to sum to 1, resulting in a trade where the highest possible score is 1.00. Anything less than ideal performance results in a lower score.

Although straightforward and easy to understand and implement, this approach suffers from one particular drawback – it is possible for an unacceptable option to score highly, provided it can overcome deficiencies in some areas with outstanding performance in others. For example, depending on the chosen importance weights, it may be possible for a powerful engine with unacceptably high cost to “win” a trade study by compensating for the poor cost score (potentially zero) with good scores in horsepower and torque. Although it is sometimes possible to manually exclude such options, this becomes more difficult if, as is often the case, evaluation criteria change during the course of the trade, or in subsequent iterations. This could make options shift in and out of the realm of acceptability, and keeping track of this can be challenging. This paper proposes an alternative method of evaluating the overall score of different options which eliminates the problem discussed above. In addition, it proposes to use importance weights in an intuitive manner which simplifies the overall process.

2. OVERVIEW OF ALTERNATIVE METHOD

This paper proposes that rather than using a weighted arithmetic sum, we use a multiplicative score based on the geometric mean, so the overall score for option n becomes:

$$S_n = \left[\prod_{k=1}^K U_{kn} \right]^{1/K} \quad (2)$$

This approach eliminates the problem described above, where an unacceptable option can still score quite well. In this approach, any option that has unacceptable performance on any criterion will score zero overall, and any option with very poor performance on any criterion will suffer a severe penalty relative to other options, and is unlikely to fare well.

Note that the multiplicative formula shown above does not include any importance weighting for the various criteria. So, for example, there is nothing in this formula to indicate that horsepower is more important than risk, or vice-versa. In order to include importance rankings in the overall score, we use the relative importance of each criterion to generate the utility curve on which it is evaluated, rather than weighting each criterion itself. This makes it somewhat easier to conduct the trade, since the user needn't create explicit utility curves for each criteria. Also, the nature of the relationship between the importance of any given criterion and the associated utility curve makes intuitive sense.

The steps associated with this trade study are similar to those for an arithmetic trade, with the exception of utility curve generation. They include the following:

1. Identify key performance and burden criteria
2. Rank the importance of these evaluation criteria
3. Set goals and limits on these evaluation criteria
4. Identify a list of options to select from, together with each option's performance on the criteria
5. Calculate scores for each option on each criterion, and calculate an overall score for each option

- Compare options head-to-head and conduct sensitivity analyses to determine *why* any option should be selected over any other

The following sections go through the first four steps in order, with examples from the engine trade study discussed above as illustration. The final two steps – score calculation and sensitivity analysis – are covered in subsequent sections.

2.1 Performance and Burdens

The first step in a trade study is to determine what characteristics we want to trade. This involves identifying both the performance criteria we desire, as well as the burdens we need to carry in order to achieve this performance. For the surrogate engine trade described here, we have three performance and three burden categories, as shown in the table below, together with the units used to characterize them. These units can be of any form – physical quantities such as time or mass, or scales ranging from 1 to 5, although numerical units are required (no categorical measures). Normally, performance criteria are “higher is better”, while burdens are “lower is better”, but there are exceptions.

Table 1: Performance and burden criteria and units.

Criterion	Units
Horsepower	hp
Maximum Torque	ft-lbs
Redline Speed	rpm
Weight	kg
Cost	\$
Risk	1-5

2.2 Importance Ranking

Once the criteria are identified, it is necessary to rank their importance. For the sake of simplicity, we use a scale of zero to ten, with zero indicating a criterion with no importance, while ten indicates very high importance. The rankings are not exclusive, so it is perfectly fine to set three of them to 5, two of them to 8, and one of them to 3. The importance rankings only serve to create utility functions that effectively weight the effect of a given criterion – the overall trade study score is more sensitive to those criteria with high

importance. The table below shows the importance rankings for this sample trade study.

Table 2: Evaluation criteria and relative importance

Criterion	Importance
Horsepower	8
Maximum Torque	3
Redline Speed	5
Weight	10
Cost	3
Risk	5

2.3 Goals and Limits

Once we’ve decided on which criteria to evaluate, we need to determine what level of performance is acceptable for each. There are three general cases – higher is better (e.g., horsepower), lower is better (e.g., cost), and nominal is best (e.g., a radio frequency). In all of these cases, we need to identify two (or three) numbers to characterize the desired performance. One is the unacceptable limit – in the event the criterion is “higher is better”, this is the minimum lower limit. Any option that doesn’t meet this minimal level of performance will score a zero in the trade study. We might know, for example, that we’ll need at least 85 hp to meet our minimal mobility and power needs. The second number is the reasonable best performance that we would like to have. For example, in an engine we would always like more horsepower, but we might know that even including growth and margin we would not need more than 150 hp, which would then be the upper limit of performance. Any engine delivering more horsepower would still score 1.0 – no added bonus for exceeding the limit. The table below shows these values for the engine trade study.

Table 3: Performance and burden goals and limits.

Criterion	Type	Best Goal	Limit
Horsepower	Higher	150 hp	85 hp
Maximum Torque	Higher	140 ft-lbs	80 ft-lbs
Redline Speed	Higher	4500 rpm	7000 rpm
Weight	Lower	100 kg	160 kg
Cost	Lower	\$2,100	\$3,000
Risk	Lower	1	5

A third number is only required for a “nominal is best” criterion, in which case we need to specify both upper and lower unacceptable limits, as well as the “reasonable best” goal.

2.4 Options and Performance

The next step in the trade is to identify potential options from which to choose. The table below shows four engine options, together with the key performance and burden data for each. For each engine, and each criteria, we need to calculate a score that represents how well that engine performs on that criterion. This is the subject of the next section – utility curves.

Table 4: Performance data for four hypothetical engines.

Criteria	Engine A	Engine B	Engine C	Engine D
Horsepower	90	120	150	125
Torque	95	90	110	130
Speed	4800	6200	6800	6500
Weight	110	130	125	150
Cost	2100	2800	2400	2300
Risk	2	1	5	3

3. SCORE CALCULATION

The calculation of scores is a two-step process. The first consists of changing the raw data numbers on individual criteria (horsepower, cost, etc.) into normalized utility scores ranging from zero to one. The second step involves compiling individual utilities into a single overall score for each option.

3.1 Individual Criteria Scores

One of the key steps in conducting the trade is constructing utility curves, whose purpose is to convert raw performance numbers (horsepower, risk, etc.) into a score between zero and one on each criterion. The figure below shows a range of possible utility curves for the horsepower criterion, where unacceptable performance is anything less than 85 hp and reasonable best performance is 150 hp.

The topmost curve in the figure shows what we would expect to see if horsepower wasn't really that important. That is, any option that has better than barely acceptable horsepower will receive a score close to one. This indicates that the analyst

doesn't care much about this measure, as long as it's somewhere in the ballpark.

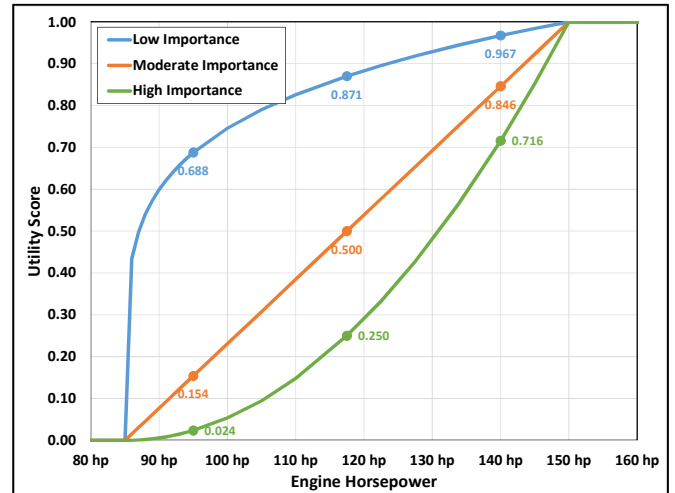


Figure 1: Utility curves for low, moderate, and highly important criteria

Conversely, the lowermost curve in the figure shows what we would expect to see if horsepower was quite important. In this case, utility falls rapidly as horsepower drops from the optimum value. So an option at the high end of the range (140 hp) scores almost 3 times better than one in the middle of the range (117.5 hp), and almost 30 times better than one near the bottom of the range (95 hp), indicating the importance of the criterion.

The middle curve in the figure – the straight line – indicates a moderate importance level. In this case, the high end performer is 1.7 times better than the average performer, and 5.5 times better than the poor performer. In all cases, once performance drops below the lower limit the score drops to zero. Similarly, once performance exceeds the upper limit, the score pegs at one.

There is a simple way to use this observation to create utility curves that take into account the relative importance of any given criterion. The formula below shows the nature of the utility curves that accomplish this for performance values between the upper and lower limits. (This equation shows a “higher is better” measure – similar curves can be constructed for “lower is better”

measures such as cost and risk, or for two-sided “nominal is best” measures.)

$$U_{kn} = \left(\frac{P_{kn} - LL_k}{UL_k - LL_k} \right)^{\left(\frac{M_k}{5} \right)} \quad (3)$$

In this equation, the unacceptable lower limit on this criterion is given by LL_k , while the reasonable best upper limit is given by UL_k . The actual performance value of any option on the given criterion is given by P_{kn} , and the importance of this criterion on a scale of 0 to 10 is given by M_k . Again, this equation applies for any performance values between LL_k and UL_k – outside of those limits, the utility is either zero or one.

The figure below shows these curves for eleven values of importance, ranging from 0 to 10. These behave as described above – criteria with low importance levels quickly reach nearly constant utility once they exceed a barely acceptable level of importance, while those with high importance change score quite quickly once they start to drop from the optimum level. Those with moderate importance show essentially linear utility.

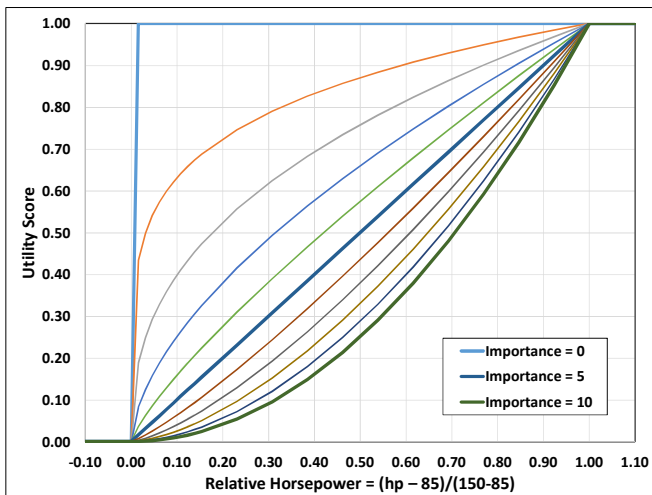


Figure 2: Utility curves from importance levels using (3).

Any criteria with zero importance have a flat utility curve that equals 1.0 everywhere within the range. This indicates that options with performance that falls anywhere within the acceptable limits are all equally desirable on this

criterion, which is consistent with an importance level of zero.

If we look at how this utility function changes with importance, we find that the slope of the utility curve is proportional to the importance level when the performance is near the high end of the range. So, the rate at which utility changes with performance is about twice as high for a criterion with an importance level of 10 as for one with an importance level of 5, and ten times that of a criterion with an importance level of 1. The same ratios hold for importance levels of 9, 3, and 1 – each changes three times more quickly than the smaller one, when the performance is near the high end of the scale. So with this function, the ratio of importance values equals the ratio of the rate at which utility changes with performance.

There is one drawback to these utility curves, and that is they are not analytic at the limits. The utility curves are continuous at 0 and 1, but the derivatives are not. There are ways to address this with more elaborate functions, and indeed it is not difficult, but for purposes of conducting trade studies the lack of continuous derivatives at the limits should have little effect. For the sake of simplicity, Equation 3 should suffice. The table below shows the individual scores associated with each option given the importance levels, goals and limits, and performance data shown previously.

Table 5: Individual scores for sample trade study.

Criterion	Imp.	Eng. A	Eng. B	Eng. C	Eng. D
HP	8	0.017	0.371	1.000	0.460
Torque	3	0.435	0.341	0.660	0.896
Speed	5	0.120	0.680	0.920	0.800
Weight	10	0.694	0.250	0.340	0.028
Cost	3	1.000	0.406	0.784	0.860
Risk	5	0.750	1.000	0.000	0.500

3.1 Overall Option Scores

As discussed earlier, the overall score for a given option is calculated as the geometric mean of the option’s individual scores on all criteria, as given in Equation (2) above. For the engine example discussed here, we would have the results shown in Table 6 and Figure 3 for overall scores.

Table 6: Overall multiplicative scores for each option.

	Eng. A	Eng. B	Eng. C	Eng. D
Score	27.7%	45.4%	0.0%	39.7%

The results show that Engine C gets an overall score of zero, despite having the highest horsepower and redline speed, and high torque and low cost. The key factor here is risk – as shown in Table 3, a risk level of 5 is considered unacceptable. As a consequence, the overall score for this option is zero – any key parameter with unacceptable performance means that the option itself is, strictly speaking, unacceptable. This is a key feature of the multiplicative nature of the method.

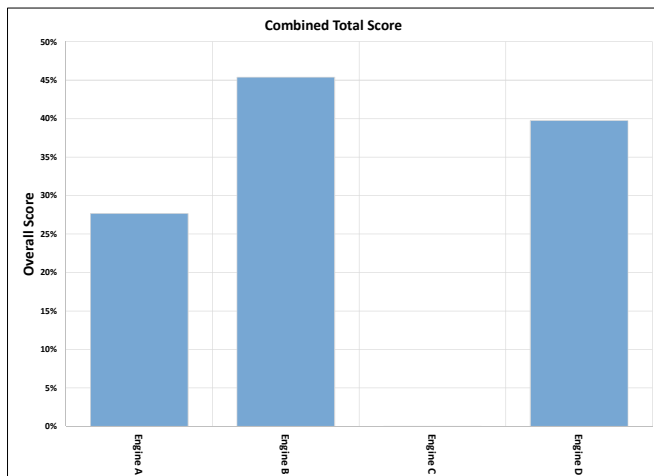


Figure 3: Multiplicative scores – hypothetical engine trade.

Of course, using an additive model, Engine C could be rejected prior to the analysis on the basis of unacceptable risk and simply left out of the trade. An advantage of the multiplicative method is that this unacceptable option can be left in, and the analysis will show that it has been considered and why it was rejected. Also, it is frequently the case that either requirements or the performance of a given option change during the course of development. Leaving an initially unacceptable option in the trade study allows it to transition into acceptability. In contrast, any option that was initially excluded from an additive trade study on the grounds of unacceptable performance on a

given criterion may fail to be added back into the trade if performance or requirements change.

In the more traditional trade study approach, the overall score is calculated as a weighted sum of the performance values. If we calculate weighting factors in a common way (individual importance divided by the sum of all importance values), the scores for the different options are very different from those shown in Table 6. We end up with the results shown below, where Engine C becomes the highest scoring option due to its horsepower, speed, torque, and cost, despite the fact that it has unacceptably high risk. The other three options are all essentially tied at 46-47%.

Table 7: Overall weighted average scores for each option.

	Eng. A	Eng. B	Eng. C	Eng. D
Score	46.3%	47.4%	59.8%	46.3%

4. SENSITIVITY ANALYSIS

Although Figure 3 shows that Engine B appears to be the preferred option, it's important to keep in mind that any trade study tool is really a decision support tool. It is not meant to make a choice of the “best” option, but to inform the analyst of why it might make sense to choose one option over another. As a consequence, perhaps the most important part of a trade study is not the calculation of overall scores, but the sensitivity analysis that compares options to each other.

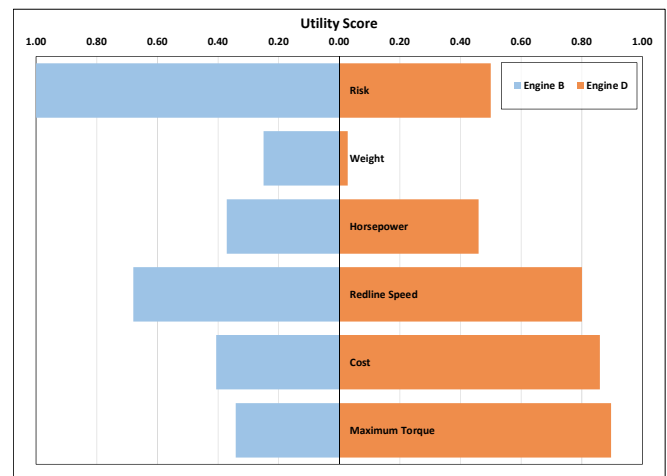


Figure 4: Head-to-head comparison of Engines B and D.

As an example, we might wish to compare the two highest-performing engines to see why Engine B scored somewhat better than Engine D. One way to do this is with a simple head-to-head comparison, showing the utility scores of each option on each measure. This is shown in two different ways in Figure 4 and Figure 5. The first shows a bar chart indicating the scores of the two measures against each other, with the length of the bar indicating the utility score. The second shows the same information in a different format (side-by-side instead of head-to-head), along with actual performance values as labels on the columns indicating utility scores.

These charts make it easy to see why Engine B is preferred to Engine D. Although B scores worse on power, speed, cost, and torque, it is somewhat better on risk and much better on weight. While the difference in weight is only 20 kg, the fact is that Engine D is within 10 kg of being unacceptably heavy, and weight is the most important criterion in the trade. As a consequence, although both options score relatively poorly in weight, Engine B scores about 9 times better.

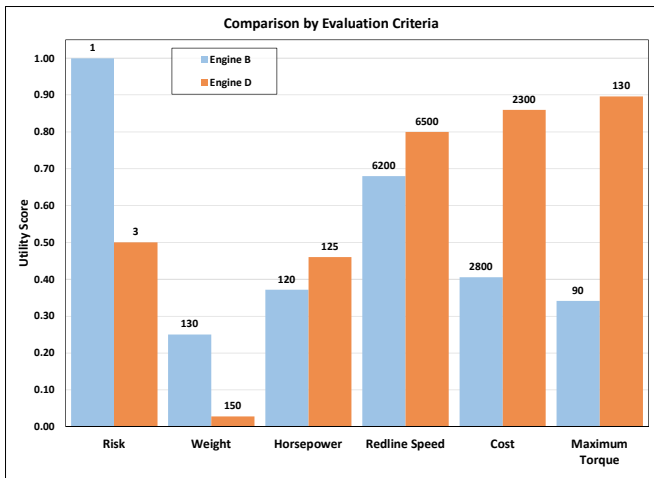


Figure 5: Side-by-side comparison of Engines B and D.

Another way to see this is to look at the sensitivity chart in Figure 6. This shows which values of importance could potentially flip the rankings of the two options. So if weight were less important (6.3 rather than 10), or if cost or torque

were somewhat more important, then Engine D would score slightly better than Engine B overall.

This indicates that the choice of Engine B over Engine D is not very robust. Relatively minor changes in importance levels can result in a change of ranking. Also, a reasonable analyst might ask himself if the 20 kg difference in weight between the two engines is enough to drive the decision. Maybe it's actually worth 20 kg in order to get better power, speed, and torque at a lower cost. Again, that's the nature of a decision support tool – it is not meant to make design decisions, but to help the analyst see the reason for selecting one option over another.

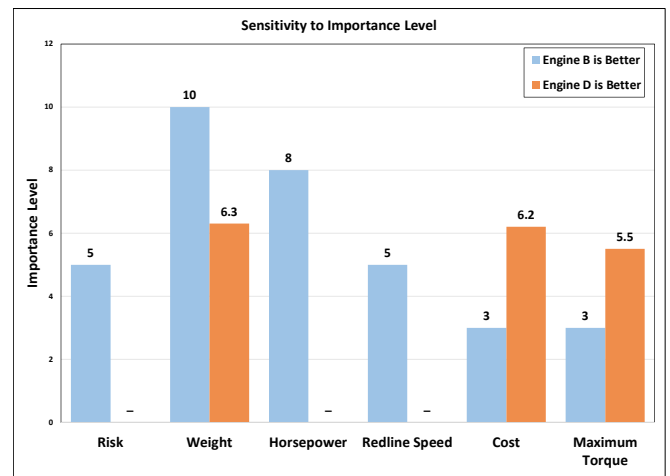


Figure 6: Sensitivity to importance levels of B and D.

Alternatively, a head-to-head comparison of Engine A with either of the other two options shows that it is difficult to find a situation in which Engine A is preferred. Only if the importance of horsepower drops to nearly zero will Engine A be preferred over Engine B. In order to prefer Engine A to Engine D, the importance of horsepower needs to drop below 3. No other factors, on their own, are capable of making Engine A the preferred option. As a consequence, the decision against Engine A seems fairly robust.

Because the unacceptable option C has been left in the trade, it is possible to include this option in any sensitivity analysis. Clearly, if risk were not a

factor, this option would no longer be unacceptable. A simple way to evaluate this is to set the importance of risk to zero. This effectively eliminates it as a criterion and allows a comparison on all the other factors. This is another advantage of leaving an initially unacceptable option in the trade study, as opposed to manually excluding it.

The sensitivity analysis shown here can also be performed on the data itself, rather than on importance values. For example, there may be some uncertainty in how much actual horsepower any option can actually provide. It is a simple matter to change data values to see how this affects the results. More generally, because the utility curves and overall score functions are analytic between the upper and lower limits, it is possible to explicitly calculate the sensitivity of any option's score to a change in any of its performance values. This is just the slope of the geometric mean function, which is the partial derivative of the geometric mean with respect to the utility of a given parameter. This is easily calculated using the chain rule for derivatives. For the utility function shown in Equation (3) above, the rate of change of the score of option n (S_n) with a variation in performance parameter k (P_{kn}) is simply:

$$\frac{dS_n}{dP_{kn}} = \frac{M_k S_n}{5K(P_{kn} - LL_k)} \quad (4)$$

In this equation, M_k is the importance of criterion k , while LL_k is the unacceptable limit of that criterion, and K is the total number of criteria.

This formula shows that sensitivity to any change in performance is directly proportional to the importance of that criterion, M_k , which is what we'd like in an importance factor. Also, owing to the $P_{kn} - LL_k$ term in the denominator, the closer the performance is to unacceptable, the more sensitive the overall score is to a change in performance. Again, this makes sense in that the difference between 0 and slightly more than 0 is very large when taking a geometric mean – it is

the difference between unacceptable and acceptable. In contrast, in a standard weighted average trade study, this difference can be vanishingly small. This analysis provides a good idea of how sensitive the results are to uncertainty in the data, which is another key factor in making design decisions.

5. EXAMPLE

The method described above is mathematically straightforward. It is even simpler to implement than a traditional arithmetic weighted sum trade owing to the fact that there is no need for the analyst to manually create utility curves. When using this method, the utility curves are a function of the importance levels as shown in Equation (3). The entire process can be programmed in an Excel workbook. The screenshots below show the menu structure from such a workbook containing the hypothetical engine trade described earlier.

Figure 7: Form to enter evaluation criteria (performance and burdens) for the hypothetical engine trade.

The Evaluation Criteria form asks for a text description of each performance goal or burden, along with the units used for measurement. These can be added or deleted with the click of a button.

Criteria	Units	Imp
Horsepower	hp	
Maximum Torque	ft-lbs	
Redline Speed	rpm	
Weight	kg	

Figure 8: Form to enter importance levels for each criterion.

The Importance Levels are straightforward and range from 0 to 10. Entering a value of zero means the criterion remains in the trade but has no effect on the scores. This is sometimes useful to examine the effect of a single criterion.

Criteria	Units	Import	Engine A	Engine B	Engine C	Engine D
Horsepower	hp	8	90	120	150	125
Maximum Torque	ft-lbs	3	95	90	110	131
Redline Speed	rpm	5	4800	6200	6800	6500
Weight	kg	10	110	130	125	150

Figure 9: Form to enter engine data for multiple options.

The Options Information form lists all the relevant data that describes the various options we wish to examine. Again, options can be added or deleted with the click of a button. Because this form sometimes contains a great deal of data, there are “zoom” buttons provided to see either all of the data or just a subset.

The Utility Information form sets the goals for each criterion. The buttons indicate whether this is a “lower is better”, “higher is better”, or “nominal is best” criterion. The form displays the minimum and maximum value of each criterion over all of

the available options, and the user then sets the goal values and the unacceptable limits. The limits can be set within or outside the maximum or minimum values – any option with performance outside the acceptable limit on any criterion will score zero.

Criteria	Units	Goal			Minimum over Options	Unacceptable Lower Limit	Target Value (Best)	Unacceptable Upper Limit
		Low	Nom	High				
Horsepower	hp				90	85	150	
Maximum Torque	ft-lbs				90	80	140	
Redline Speed	rpm				4800	4500	7000	
Weight	kg				110	-	100	

Figure 10: Form to enter goals and limits for each evaluation criterion.

Once all the data is entered, the tool displays the results of the trade along with a menu that allows the user to conduct sensitivity analyses, look at utility curves, plot head-to-head comparisons, and change importance levels. This is all part of the sensitivity analysis that is crucial to understanding the motivation for selecting one option over another.

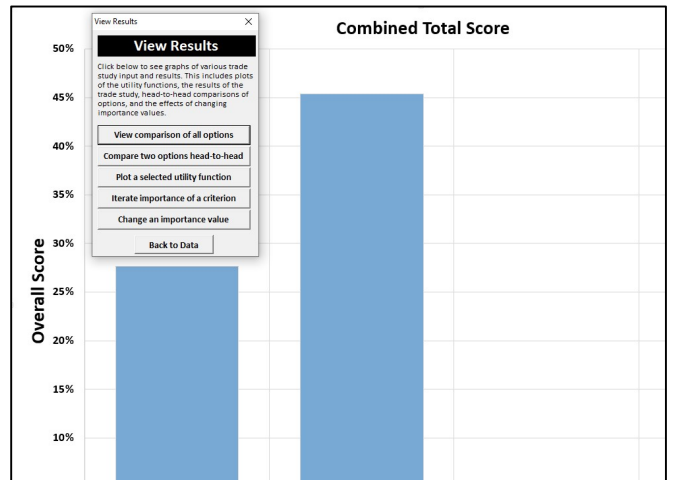


Figure 11: Overall scores, and menu to select additional options such as head-to-head comparisons.

The advantages of setting the method in Excel include the fact that it is widely available standard

desktop software that requires no special license. In addition, it is straightforward to program and simple to use. The charts and data generated by the trade are already in a workbook and can be tailored for any purpose, such as formatted for a designated presentation style. The workbook itself is intuitive, and the data from the trade can be easily exported in .csv file format for archive purposes. Typically, all the data that characterizes a trade can be held on one worksheet with a file size of a few kilobytes, and as a .csv file it is readable as text.

6. CONCLUSION

This approach to trade studies can support design decisions at the system and sub-system level, quickly identifying the most likely options, providing head-to-head comparisons between any two, and generating sensitivity charts that show the robustness (or lack thereof) of the choice. It is easy to implement and can be coded in a workbook that is intuitive to learn and utilize. It is somewhat simpler than conventional weighted

sum methods in that utility curves are automatically generated from importance levels, eliminating a step in the process. The results of the study make it easy to distinguish which options have unacceptable (or nearly unacceptable) performance on any key criteria, which is one of the significant advantages of this method over the conventional weighted sum approach.

7. REFERENCES

- [1] David Strimling, "Use of Multiple Criteria Decision Analysis in the Marine Corps Advanced Amphibious Assault Vehicle (AAAV) Program", 67th MORS Symposium, 1999.
- [2] James Gerlach, et.al. "Achieving Balanced Systems through Structured Decisions", 2013 GVSETS, Novi, Michigan, 2013.
- [3] A. Kossiakoff and W.N. Sweet, "Systems Engineering Principles and Practice," 2003.
- [4] <https://www.cbc.devcom.army.mil/train-and-support/decision-analysis/analytical-tools/>
- [5] <http://www.logicaldecisionsshop.com/catalog>