




WE'VE GOT TWO
STAR WARS PLANS,
MR. PRESIDENT.




PLAN A STOPS 80%
OF ALL SOVIET
MISSILES...




OF COURSE, THE OTHER
20% WILL GET THROUGH...



PLAN B HAS AN 80%
CHANCE OF STOPPING
ALL SOVIET MISSILES...



BUT A 20% CHANCE
OF **EVERYTHING**
GETTING THROUGH.



WHAT
IS
YOUR
DECISION?



NOTES ON MEASURES OF EFFECTIVENESS

Naval Postgraduate School
Monterey, California

NOTES ON
MEASURES OF EFFECTIVENESS

8/19/85

by
Edward B. Rockower
Dept. of Operations Research
Naval Postgraduate School
Monterey, CA. 93943

INTRODUCTION

[*"Man is the measure of all things"*, Protagoras]

We begin our discussion of Measures of Effectiveness (MOE's) with a description of the context within which they're defined. If Operations Research is "a scientific approach to problem solving for executive management"¹ then it's clear that the focus is on executive decision making, as well as on the techniques for solving problems. Operations Research involves constructing mathematical, economic, and statistical descriptions or models of decision and control problems in order to treat situations of complexity and uncertainty.² In general, to apply Operations Research to any field, we must:

1. define the system we are dealing with (e.g. the level of detail or aggregation, what parameters are exogenous and which are endogenous, which are decision variables and which are dynamical or intermediate variables, which are deterministic and which random variables, etc.)
2. define an MOE which tells us how well we are doing in making decisions, and how well pleased we are with the outcomes resulting from our actions. This lets us rank the outcomes resulting from alternative courses of action (or strategies).
3. construct a model, whether analytical, verbal, computer simulation, or whatever, which allows us to predict the results of our decisions.
4. gather a data base which contains those exogenous variables, system parameters, physical constants, etc. which must be fed to our model, along with values of our decision variables.
5. optimize our solution; i.e. determine those values of our decision variables within the feasible region (or domain) which result in the maximum values of our MOE.
6. report, explain, communicate, and interpret so that the qualitative, judgmental, political and operational impacts can be assessed by the ultimate decision maker... this must be successful so that any additional, non-quantifiable factors can be combined, by the decision maker, with the analysis. This also aids in the implementation stage, so that those tasked with the job of implementing the decision will understand and back it.

¹ Harvey Wagner, Principles of Operations Research, 2nd Edition, Prentice Hall, (1975)

² *ibid*, cf also Naval Operations Analysis, The Naval Institute Press (1978)

In other words, in order to construct methods to aid us in making better decisions, one of the first things we must do is establish a consistent, quantitative, measurable, and credible, measure of 'how well we are doing' in trying to achieve our goals. We must have a means of assessing the value of alternative courses of action to the decision maker. There are a number of related terms, some of which are synonymous with MOE, e.g.

- a) measure of performance (MOP)
- b) index of effectiveness³
- c) figure of merit (FOM)
- d) operational effectiveness
- e) value, utility, cost ...
- f) benefit - cost ratio (B/C or cost/benefit ratio)


We'll define and apply some of these but it should be remembered that the definitions and usage of these terms is not uniform or standardized in practice. In general there exists a hierarchy of MOE's or FOM's. What serves as an engineering measure of performance at one level of the hierarchy (e.g. rate of fire of a gun), may be viewed as the MOE at a lower level. Similarly, at an intermediate level of the hierarchy operational effectiveness may be viewed as the MOE in an evaluation of tactics but at a higher level only the Pr[win the war] will be the appropriate MOE. In fact, at each level there is certainly at least one "measure" which indicates how well you're doing, from the narrow perspective of a local objective, e.g.

- a) E(no. attackers required to destroy a target)
- b) Pr[target is destroyed]
- c) E(no. attack aircraft destroyed per SAM available)
-
- d) Pr[win the battle]
- e) E(time, T, until battle is won)
- d) Pr[T≥34 hours]
-
- e) Pr[win the war]
- f) E(cost to win war)

³Venntsel, Ye. S., Introduction to Operations Research, Soviet Radio Publishing House, Moscow (1964)

An example⁴ of the hierarchy of effectiveness may be visualized as the following:

FACTORS USED AS MOE



We will return to this matter of selecting the measure appropriate to our position in the hierarchy after developing some machinery with which to deal with MOE's.

Whichever MOE we select, in each case we wish to solve the mathematical model for the optimal value of the decision variables (represented by the vector $\underline{x}$) i.e. those values which result in the maximum (or possibly, minimum) value of the MOE. Let $F(\underline{x})$ represent our (possibly vector valued) model of the operational system under study. Then, stated mathematically, our problem is to find

$$\underline{x}^* = \arg. \text{MAX}_{\underline{x}} F(\underline{x}).$$

Of course, there may be a number of output MOE's, since $F(\cdot)$ may be vector valued, e.g. $F(\cdot)$ may generate $\text{Pr}[\text{win the battle}]$, $E(\text{no. of casualties})$, cost in time to win the battle, etc. Often, we want to combine these MOE's into one overall MOE by taking some kind of weighted average of each individual MOE. Unfortunately, there are problems of scale and units. Probabilities are between 0-1 and have no units, time may be in units of minutes, hours, or days, and may be any positive number

⁴ U.S. Army TRADOC Study

from 0-thousands, human life is difficult if not impossible to quantify, and in any case, is incommensurable with minutes, etc. Also there are uncertainties regarding how to weight the relative importance of each factor in any overall MOE. Often the appropriate solution is to present these components of MOE to the ultimate decision maker (DM) and allow him to combine them based on his experience and judgement. It is well to remember that neither the analyst nor his model makes the decisions, they just calculate the consequences of alternative decisions.

We can summarize this introduction by stating the basic modelling questions from another perspective:

- a) is a decision required
- b) what do we need to decide? (i.e. what is the Domain of $F(\underline{x})$ or what is the list of options available to the DM)
- c) what are we trying to do? (i.e. what is the Range of $F(\underline{x})$ or what MOE do we want to maximize)
- d) how does a decision affect the MOE? (i.e. what is the function $F(\underline{x})$ or how do we specify our model of the operational system)
- and , finally,
- e) what decision optimizes the MOE? (i.e. how do we mathematically solve the optimization problem; numerous techniques of Operations Research are used here; e.g. Linear Programming, Non-Linear Programming, Dynamic Programming, Critical Path Method, Game Theory, computer simulation, exhaustive enumeration...)

From the above we see that a crucial part of the initial analysis⁵ is selection of the appropriate measure of effectiveness. We take up the details of this problem in the rest of these notes.

PROPERTIES OF MOE'S

We can almost immediately identify a number of properties that a reasonable MOE must possess:

- 1) a measure of effectiveness must be closely related to the objective of the operation. In other words, it should serve as a good surrogate for the real goals. For example, the number of submarines sunk per month may be OK as an MOE if the real objective is to destroy submarines, but if the objective

⁵ also see Raisbeck, Gordon; "How the Choice of Measures of Effectiveness Constrains Operational Analysis", Interfaces, 9, no. 4, August 1979

is to protect shipping then the best actions may imply sinking fewer subs, e.g. a convoy may be better served with alternatives which result in evading the subs more effectively.⁶

2) an MOE must be measurable and quantifiable, and the data with which to calculate it must be available, or such that one can obtain it. Then the MOE can be calculated, and used to form a basis from which to make decisions.

Example #1: Sonobuoy selection

In order to evaluate a number of alternative available sonobuoys there are a number of possible MOE's, for example,

- 1) $MOE_1 = \text{radius of coverage} = R$
- 2) $MOE_2 = \text{sweep width} = 2R$
- 3) $MOE_3 = \text{coverage area} = \pi R^2$
- 4) $MOE_4 = \text{Pr}[\text{sinking target using sonobuoy} \mid \text{given value of } R]$

The operations analyst must select one of these, or possibly other, MOE's which most closely represents the stated goals of the executive decision maker.

The selection of MOE is made somewhat easier by the existence of the so-called "scaling rule" which holds so long as the outcomes resulting from a decision are deterministic and ratios of benefit/cost are not considered.

The Scaling Rule: If $MOE_a = h(MOE_b)$ where $h(\cdot)$ is any monotonically increasing function, then one says that the two MOE's are "Decision Equivalent" (DE), i.e. MOE_a and MOE_b will rank the alternative courses of action (the alternate strategies or values of the decision variables) in exactly the same order of desirability.

i.e. if $MOE_a(R) < MOE_a(R')$, then also
 $MOE_b(R) < MOE_b(R')$, and vice versa.

In our example, the first three MOE's are decision equivalent, and in fact, if $MOE_4(R)$ is monotonically increasing with R then we don't need to know its exact functional form, it is DE with the other MOE's (although that's not likely to be true, i.e. a large R may be worse than a smaller value of R unless the sonobuoy also gives us information about range and direction; if $R = 5000$ nm then the sonobuoy may just tell us that the sub is within our ocean!).

⁶Naval Operations Analysis, p.13

Now, what if the model output is not a single scalar or vector quantity but instead is a random variable? Then each decision option x can result in many different outcomes. Our model will (most generally) generate a probability distribution (in the discrete outcomes case) or a probability density function (pdf, in the continuous case). Our MOE must then rank probability distributions over values, instead of a deterministic scalar (or even vector).

For example, what if R for our sonobuoy depends on ambient acoustic conditions in the ocean or on an imperfect reliability of the sonobuoy. Then our two options might become:

sonobuoy #1 $R = 0$ nm, with probability .5
 $R = 4$ nm, with probability .5

sonobuoy #2 $R = 2.5$ nm, with probability 1.

Which is preferred? Should we take the average value of R_1 as our MOE, where $E(R) = 2$? Comparing this with $R_2 = 2.5$ we would prefer sonobuoy #2.

On the other hand, should we use $E(\text{area coverage})$?

$$E(\pi R_1^2) = 8\pi$$

$$E(\pi R_2^2) = 6.25\pi$$

with this as our MOE we would prefer sonobuoy #1. Hence the rankings are reversed, depending on which MOE we use, so they are no longer decision equivalent. To confuse matters further, it may be that the task requires that the sonobuoy have a range of at least 2 nm. Sonobuoy #1 will perform adequately only 50% of the time, whereas #2 meets the requirement 100% of the time. Hence, according to the MOE, $\Pr[R \geq 2]$, we would once again prefer sonobuoy #2.

PROBLEM: One of the above two sonobuoys must be selected to be used in a barrier across an important channel. If detection of an enemy submarine occurs no less than 3nm upstream of the barrier there will be sufficient time for a P-3 to respond and take appropriate action before the enemy submarine passes the barrier and loses itself in the ocean.

- a) what is the appropriate MOE?
- b) which sonobuoy should be selected?

Most generally, it is really the entire pdf which must be used to rank outcomes, not just its mean, cumulative probability, variance, etc.

The above example has demonstrated that

1) even when 2 or more MOE's are decision equivalent with respect to "certain" outcomes (i.e. the outcome resulting from each choice is known with certainty, or prob = 1) they do not give the same ranking of alternative courses of action when the outcomes must be described by probability distributions.

and 2) although the natural thing to do seemed to be to use the expected value of each MOE as a way of taking into consideration the pdf, its clear after a little thought that not only does that not lead to consistent rankings of alternatives but in fact a CDF may be more appropriate at times, or other parameters of the distribution.

UTILITY THEORY

["Ah, if the rich were rich as the poor fancy riches!", Emerson]

In order to develop a means to rank probability distributions we must resort to what is called the theory of "Utility". This provides us with a consistent way in which to combine different figures of merit or sub-MOE's, subjective and judgemental factors, and probabilistic aspects of outcomes (i.e. risk factors) into an overall MOE appropriate for the level in the hierarchy in which each decision maker must function.

FIRST EXAMPLE: Buying a Lottery Ticket

'Dollars' is probably the MOE you think you're using when trying to decide whether or not to buy a (State) lottery ticket. However, a little calculation for a typical lottery will quickly show that the expected return for a dollar invested is on the order of 50 cents. In other words, there is an expected loss of about 50 cents resulting from a decision to invest one dollar in purchasing a lottery ticket. Let's frame the decision in terms of the following two alternative courses of action (let's call them both "lotteries", by definition):

<u>choice</u>	<u>outcome</u>	<u>prob.</u>
#1 (don't buy the ticket)	keep \$1.00,	1.
#2 (buy lottery ticket)	\$0.00 (lose),	1-10-6
	\$500k(win),	10-6

The expected payoff (or expected monetary value EMV) of the second lottery is easily calculated,

$$E(\text{payoff}) = 0 \cdot (1 - 10^{-6}) + \$500,000 \cdot (10^{-6}) \\ = \$0.50$$

Hence, one might think that a rational decision maker would look at the above two alternative courses of action,

1.) - invest in "lottery" #1 or 2.) - invest in lottery #2,

note that the expected payoff for the first lottery is larger than for the second, and decide to invest in the first one. In other words he won't buy the state lottery ticket.

That rational people (not just compulsive gamblers) do, in fact buy lottery tickets, in spite of the fact that the expected payoff is less than for keeping their money, or investing in something else, challenges us to try to determine the actual MOE they are using in making their decisions. Whatever the MOE they are using, clearly it ranks the above two alternatives in the opposite order to the ranking given them by $E(\text{payoff})$.

SECOND EXAMPLE: Double or Nothing on your Salary

Which of the following two gambles would you prefer?

	<u>choice</u>	<u>outcome</u>	<u>prob.</u>
#1	double salary (one year's)	----- +1 salary	.5
		----- (-) 1 salary	.5
#2	double \$100.	----- +\$100.	.5
		----- (-) \$100	.5

Let's assume that the salary is net, after taxes and deductions. It's easy to see that the expected payoff from both lotteries is 0. Still, most people would rather bet double or nothing on \$100. than on next year's salary. This isn't too hard to understand. They'd like to avoid the risk of losing an entire year's salary. Even if \$50 were added to each of the prizes in #1 you'd probably still prefer #2. In other words, the "real" subjective loss from losing a year's salary is generally far greater in magnitude than the real gain from receiving double the usual amount. Hence, we might say that the average decision maker is "risk averse" with respect to this decision. In our previous example, by choosing to invest in the lottery the DM shows he values more highly the one-in-a-million chance of winning \$500,000 than the very likely loss of \$1.00. We might say he is "risk preferring" with respect to that decision.

But note, there may be circumstances in which almost everyone would prefer the first alternative "lottery" in our last example. For example, what if a close relative needed an operation and this lottery is the only way we have a chance of obtaining the needed money to pay for it. In that case we stand to gain more by winning the first lottery (where we may gain an additional year's salary and thereby save the life of someone close to us) than we stand to lose. So, given the right circumstances, any one of us would be likely to change from risk averse to risk preferring. Same person, same lottery, only our assessment of the relative value of each of the lotteries has reversed.


Finally, consider the following historically important problem, known as the St. Petersburg Paradox. How much would you be willing to pay for the right to play the following game? A fair coin is flipped until it first comes up tails. If the first appearance of tails is on the n^{th} toss then you'll receive 2^n dollars. Before reading ahead, write down the number of dollars you'd be willing to pay to play this game. \$_____

PROBLEM #1: Show that the probability of tails appearing for the first time on the n^{th} toss is given by a geometric distribution. Calculate the expected value of n , $E(n)$, and the expected payoff. Show that the latter is infinite.

As you show in problem #1 the expected monetary value to the person lucky enough to play this lottery is infinitely large. Hence one might think that a rational decision maker would be willing to pay a very large amount of money to play this game. Certainly all the money you have in the bank can't be too much to pay. But you probably didn't write down much more than \$20 to \$50, if that much. In fact, it was recognized hundreds of years ago that people were simply not willing to pay very much money to play this game. To explain this apparent paradox, Daniel Bernoulli reasoned that the real value, or "utility", of an additional increment of money depends on how much you already have. Clearly a millionaire who receives \$10,000 would derive far less emotional satisfaction, and his life would be improved to a lesser extent, than would a person who earns \$10,000 per year. In fact, Bernoulli assumed that the additional value (or utility) of an added increment of money was inversely proportional to the amount of money one already possesses.

i.e. $dU = c \cdot d\$/\$$,


where c is some proportionality constant. This is easily integrated, yielding the result that utility is proportional to the logarithm of the total quantity of money one possesses. Now, the log function is concave downward, consistent with the idea of diminishing additional utility as one's wealth increases, from each increment of added money. The figure shows the general form of utility vs. dollars for this type of utility function.



The specific functional form of Bernoulli's utility function should not be taken too seriously. However, the general ideas he proposed do explain qualitatively the risk aversion we often see in people's choices, for instance in the unwillingness of most people to risk next year's salary in an apparently fair lottery.

How does one explain situations in which decision makers show risk preferring behavior, as in buying a state lottery ticket? We must look for utility functions which also have concave-upward segments. Again, this can be understood on intuitive grounds. While it may be true that additional increments of money have diminishing additional utility for the average person (say, up to \$100,000) there is a point at which one has enough money to significantly change one's life. Maybe you feel that \$1,000,000 will make you financially independent. Hence \$1,000,000 is somehow qualitatively more than, say, \$100,000. Obtaining \$1,000,000 may propel you onto a new plateau of benefit. Perhaps you can then afford to devote your time to some currently unobtainable career or interest. The following utility curve shows the general form. Note that the concave-downward segment of the curve will result in risk-averse decisions while the concave-upward segment will result in risk-preferring

decisions. In other words, if the possible outcomes of a lottery fall entirely within the concave downward (upward) portion of the curve, then the DM will appear to be risk averse (preferring) with respect to that lottery.



Now let's consider an example from a combat situation, where perhaps the commensurability and relative values of outcomes are even more difficult to assess.

THIRD EXAMPLE: Holding a Position

As commander of your troops, you're given the order to hold a position to which there are 12 possible approaches. You start with 20 men in your command. As long as you have at least 12 men left you can guard each of the 12 approaches, but if you suffer more than 8 casualties, you must abandon the position. You have the responsibility of judging the relative importance (value or utility) of the two goals 1) maximize the number of survivors under your command and 2) maximize the probability that you will be able to hold your position. It seems reasonable that two casualties are approximately twice as bad as suffering one casualty. Perhaps three casualties are about three times as bad as only one casualty. But clearly 9 casualties (and having to abandon the position) is far worse than 9 times as bad as one casualty. In fact, looking at it in terms of survivors, there is a step increase in utility for having 12 survivors as compared with having only 11 survivors. Hence, how do you compare two possible alternative plans for carrying out a necessary mission, each with different estimated numbers of casualties. Let's say that plan #1 has a .5 probability of 10 casualties (and .5 probability of 0) while plan#2 has an 80% chance of incurring 8 casualties.

For plan #1

$$E(\text{no. casualties}) = 5.0; \quad \Pr[\text{hold pos'n}] = .50$$

while for plan #2

$$E(\text{no. casualties}) = 6.4.; \quad \Pr[\text{hold pos'n}] = 1.00$$

Hence, if the commander is willing to suffer 1.4 additional expected casualties he will be assured of being able to hold his position. Neither the expected number of casualties nor the probability of holding the position are adequate by themselves for determining the relative ranking of the two plans. Rather, it is the complete probability distribution which must be considered, along with the relative value, in the commander's judgement, of holding the position and of his men's lives. Stated differently, the commander must analyze the situation carefully enough so that he is clear as to what the appropriate MOE should be. Should he use

MOE = $a * E(\text{no. casualties}) + b * \text{Pr}[\text{hold pos'n}]$,

or should he use

$$\text{MOE} = a * E(\text{no. casualties}) * \text{Pr}[\text{hold pos'n}],$$

or what?

How then can the commander factor in both the complete probability distribution of outcomes for each plan, as well as the relative value of each outcome, so that he can consistently rank the plans?

The theory of expected utility provides "a means of encoding the decision makers risk preferences so that consistent choices can be made among alternative courses of action (strategies or plans) whose outcomes are uncertain". We will need some definitions in order to develop enough of this theory for our purposes.

A lottery is any uncertain proposition with specified prizes and probabilities of receiving each prize. Graphically, we depict a lottery as a type of probability tree, e.g.

lottery L ----- win salary, prob = .5

```
lose salary, prob = .5
```

or,

lottery L' ----- win ton of gold, prob = p_1
 ----- win ton of cheese, prob = p_2
 ----- win ton of poisonous snakes
 prob = p_3

The lottery L' could be either very favorable or not favorable at all, depending on the probabilities of receiving the three prizes, and also on how the decision maker compares the value of the prizes. It would be convenient to be able to convert them all to some common measure of value, e.g. to dollars, or to "utils" (arbitrary measure of utility), or something. When the prizes can all be measured in terms of the same quantifiable commodity or units then the lottery can be represented as a random variable (r.v.). Let's imagine this has been done. Then the expected value of the lottery (L) with prizes g_i which have probability p_i is,

$$E(L) = \sum p_i * g_i.$$

A decision maker who chooses between lotteries according to expected value, $E(L)$, is an "expected value decision maker".

Selling Price: The selling price of a lottery "L", abbreviated $SP(L)$, is the minimum amount of money (or some other agreed upon unit of value) that a seller would have to be paid to give up that lottery. This is aka the "certain equivalent" of the lottery. For example, imagine you owned the right to toss a fair coin and receive \$500 if it comes up heads and \$0 if it comes up tails. If you'd be willing to accept \$150 for that right but not \$149 then the selling price, $SP(L)$, is \$150. Clearly you'd also be willing to accept \$155, or \$300, or any amount over \$150. So, in other words you'd accept a certain \$150 vice an expected return of \$250 from the lottery. Now, if your total assets are "T" then you're indifferent between the certain lottery $T+SP(L)$ and the uncertain lottery $T+L$. In graphical form,

----- $T+SP(L)$, $p=1$ ----- $T+500$, $p=.5$
 ~
 ----- $T+0$, $p=.5$

where the twiddle, $\sim$, indicates indifference between the lotteries.

In general, $SP(L)$ is a function of an individual or organization's attitude toward risk and their total assets. For an expected value decision maker

$$SP(L) = E(L)$$

and we say that this DM is risk indifferent. If the DM has

$$SP(L) > E(L)$$

then we say he is risk preferring, and if the DM is such that

$$SP(L) < E(L)$$

then he is risk averse. Most individuals and organizations are generally risk averse with respect to most decisions. This is more or less for the same reasons that Bernoulli first suggested.

Buying Price: The buying price of a lottery, $BP(L)$, is the maximum amount of money a buyer would pay in order to own the lottery. In other words, a buyer with assets T is indifferent between the certain lottery T (i.e. keep your money and don't get to play the lottery) and the uncertain lottery $T - BP(L) + L$. Graphically, for the coin tossing lottery this would be expressed as,

$$\begin{array}{ccc} \text{----- } T, p=1 & & \text{----- } T - BP(L) + 500, p=.5 \\ & \sim & \\ & & \text{----- } T - BP(L) + 0, p=.5 \end{array}$$

For an expected value DM $BP(L) = SP(L) = E(L)$.

As we saw above, Daniel Bernoulli first introduced the quantitative notion of a utility function. Although his logarithmic type of utility function, $U(\cdot)$, could explain how risk aversion comes about, it doesn't explain risk preferring behavior. It is well to bear in mind that utility functions are probably more realistic when measured, than when derived. We will show below how a DM's utility function can be determined. First we will present 4 axioms which, if we subscribe to them, insure the existence of a utility function over lottery prizes. This utility function, $U(\cdot)$, will have the desirable property that, if L and L' are any two lotteries, then our DM will prefer L over L' if and only if (iff) the expected value of the utility of lottery L is greater than that of L' , i.e.

$$E[U(L)] > E[U(L')],$$

where

$$E[U(L)] = \sum p_i * U(g_i),$$

and the g 's are the prizes or payoffs for the lottery. So, there is an MOE that reflects our ranking of outcomes even in the presence of uncertainty. With the appropriate utility function the DM or his analyst can consistently rank all alternative courses of action (lotteries) at his disposal. He can, in fact, use $E[U(L)]$ as a consistent measure of effectiveness of a tactical plan, a sonobuoy, or a weapon

system, when more analytically derived MOE's such as $E[\text{no. casualties}]$, reliability, $E[R]$, $E[\pi R^2]$, cost, etc. are inadequate.

Axioms of Utility

1. Transitivity: For any prizes a and b , one of the following must hold

a is preferred to b ($a > b$)

b is preferred to a ($b > a$)

or, the DM is indifferent between them ($a \sim b$). The ordering must be transitive (as in arithmetic), i.e. if a is preferred to b and b is preferred to c then the DM must prefer a to c , symbolically,

$$a > b \ \& \ b > c \Rightarrow a > c.$$

Also, indifference is transitive,

$$a \sim b \ \& \ b \sim c \Rightarrow a \sim c.$$

Individuals may appear to violate this when looking at different aspects of the prizes, e.g. we may prefer the gas mileage of one car but the sportiness of another. Hence the DM must reduce each prize to a scalar measure of value, e.g. a dollar value.


2. Continuity: if $a > b > c$ then there is a probability, p , such that the DM is indifferent between receiving b with certainty or having the lottery in which he receives a with probability p and c with probability $1-p$, i.e. using a fairly obvious notation: $b \sim [p, a; (1-p), c]$.

In fact, querying the DM regarding the p at which he becomes indifferent is a way of determining his utility scale.

3. Monotonicity: if $a \succ b$ then; $[p, a; (1-p), b] \succ [p', a; (1-p'), b]$ iff $p \geq p'$.

In words, if the DM either is indifferent between the prizes a and b or prefers a to b , then he will prefer that lottery which produces the preferred prize with the greater probability.

4. Decomposability: aka the "no fun in gambling" axiom. The DM is indifferent between compound and simple lotteries as long as the probabilities of receiving each of the prizes are the same. For example,



Clearly, there is sometimes fun in gambling, and this axiom has been questioned in the recent operations research literature⁷. Nonetheless we will assume its validity here.

As stated above, if these axioms are assumed to be satisfied, then it can be proven that there exists a utility function, $U(\cdot)$ such that, if A, B are two lotteries or alternative strategies with uncertain outcomes, then

$$A \succsim B \text{ iff } E[U(A)] \geq E[U(B)],$$

so if we know the utility function (for a specific DM with respect to a specific class of decisions) then we can predict what the DM will choose, and we can help him to choose consistently among alternatives.

We now summarize a number of properties of utility functions.

1. The utility of any lottery is the expected utility of its prizes. (This tells us how to treat outcomes that are random variables).

2. The above equation is strictly true only when the prizes in lotteries A and B contain the DM's total assets. We should really write,

$A+T \succsim B+T$ iff $E[U(A+T)] \geq E[U(B+T)]$, where T is the DM's initial total assets, apart from the lotteries.

3. $U(\cdot)$ is unique only up to a positive linear transformation. In other words, multiplying U by a positive constant and adding an arbitrary constant to the definition of U in the above inequality will not change the sense of the inequality. Hence, preferences will remain unchanged under this type of utility scale change.

4. "More is not worse", i.e. $U(\cdot)$ is monotonically increasing. (If you're in a lifeboat and you have 5 tons of food and water dumped on your boat, you might question this as your boat sinks. Some reasonable constraints on "more" must be assumed.)

5. When $U(\cdot)$ is concave downward (e.g. Bernoulli's log function) then the DM is risk averse and $BP(L) < E(L)$.

6. When $U(\cdot)$ is concave upward then the DM is risk preferring and

⁷Bell, David, "Disappointment in Decision Making Under Uncertainty", Opns. Res., 33, Jan-Feb, pp. 1-27, (1985)

$$\text{BP}(\mathbb{L}) > \text{E}(\mathbb{L}).$$

Two extreme cases that can be identified are:

7. When $U(\cdot)$ is a linear function of the prizes, (e.g. $U(x)=a + b \cdot x$) then using $E[U(X)] = a + b \cdot E[X]$ as the MOE is decision equivalent to using $E[X]$. In other words, an expected value decision maker has a linear utility function in the region of interest.

8. When $U(\cdot)$ is a step function (e.g. having value "a" for $x < x_0$ and "b" for $x \geq x_0$) then using $E[U(X)] = a + (b - a) \cdot \Pr[x \geq x_0]$ as the MOE is DE to using $\Pr[x \geq x_0]$.

FOURTH EXAMPLE: Assume a DM has a utility function $U(x) = \sqrt{x}$. He has total assets of $T = \$200$, plus the rights to the following lottery,

A: \$100., prob .8

(- \$50.), prob = .2

a) What is the Selling Price, $SP(A)$, of the lottery, i.e. what the certain price SP such that he is indifferent between $\$200 + SP$ and $\$200 + A$?

To answer this, we seek SP which solves the following equality,

$$U(200 + SP) = E[U(200 + A)].$$

In terms of the square root utility function,

$$\sqrt{(200 + SP)} = .8 * \sqrt{300} + .2 * \sqrt{150}.$$

We then obtain $SP = \$65.88$. But note, the expected payoff from the lottery is given by $E(A) = .8 \cdot 100 + .2 \cdot (-50) = \70 . Hence, $SP(A) < E(A)$ which implies that the DM is risk averse. This was to be expected from the fact that the square root function is concave downward.

b) If the DM didn't have A, what would be his buying price, $BP(A)$?

We must solve for BP in the equation $U(200) = E[U(200-BP+A)]$, or,

$$\sqrt{200} = .8*\sqrt{(300-BP)} + .2*\sqrt{(150-BP)}.$$

The solution to this yields $BP(A) = \$64.07$, which is still less than $E(A)$. Can you see from the concave shape of the curve why the BP is less than the SP?

We now turn to the problem of setting up a utility scale for a DM. It must be borne in mind that the utility function derived is only for a specific type of decision (e.g. combat alternatives, personal investments, psychological benefits from alternative vacations, etc.). Also, the utility function is only valid for a specific decision maker whether an individual, institution, or nation. The function $U(.)$ will generally not be simply concave upward, downward or a straight line everywhere on its domain. Hence the DM cannot be characterized as risk averse, preferring or neutral except when his total assets and the specific alternative choices are specified. Then he may be, e.g., risk averse for this decision, given his current total assets.

Clearly in a war-fighting environment the best MOE would allow one to rank outcomes and alternative strategies with probabilistic outcomes according to their impact on winning the war (whether the prob. of winning, the $E[\text{time to win war}]$, etc.). Usually this is unrealistic. It is just too difficult to determine the quantitative impact of holding a particular position, conserving a scarce resource, or even winning a particular battle, on the outcome of the war effort. Barring availability of this kind of (almost omniscient) knowledge we must use a surrogate, intermediate MOE which we hope is positively correlated with the success of the total war effort. This local MOE will be more tractable since it will deal mainly with more local considerations than for the entire war effort. Clearly we can expect that holding a particular hill is "good" for the war effort, although it would be difficult to quantify its impact even on the probability of winning the current battle, let alone on the probability of winning the war. We must rely on the judgement and experience of the local DM to factor in more global considerations and to assess the relative importance of various outcomes and of their impact on an unspecified global MOE.

With this philosophy, and with the understanding that there are, in fact, many situations in which one can more objectively derive reasonable quantitative MOE's, we proceed to set up a utility scale for a specific DM with a hard choice to make.

Setting Up the Utility Scale

Refer to the example of holding the position for which there are 12 possible approaches and the DM initially has 20 men. Using the freedom that the utility function is arbitrary up to a positive linear transformation, for convenience we define

$$U[20 \text{ survivors, i.e. no losses}] = 100,$$

$$U[0 \text{ survivors, i.e. 20 casualties}] = 0.$$

Now, using the axiom of continuity, we know that there is a probability p such that the DM will be indifferent between obtaining outcome "i" with 100% certainty (e.g. 14 survivors from a required

mission) and the lottery in which there are 20 survivors with probability p and 0 survivors with probability $1-p$, i.e.

$$\begin{array}{ccc} \text{----- outcome } i, \text{ prob}=1 & & \text{-----20 surv, prob}=p \\ & \sim & \\ & & \text{-----0 surv., prob}=(1-p). \end{array}$$

At that point of indifference between the above two lotteries,

$$\begin{aligned} U[\text{outcome } i] &= p \cdot U[\text{no losses}] + (1-p) \cdot U[\text{no survivors}] \\ &= p \cdot 100, \text{ for the scale we've chosen.} \end{aligned}$$

We ask the DM at what value of p his indifference point is reached. He may tell us that he would prefer a certain 3 casualties to having no losses with .95 probability and 20 casualties with .05 probability. However, he answers that he is indifferent between a certain 3 casualties and the alternative of no casualties with .98 probability and all lost with probability .02. This gives us another point on our utility curve, since we now know that

$$\begin{aligned} U[17 \text{ survivors}] &= .98 \cdot 100 \\ &= 98. \end{aligned}$$


Continuing in this way, we may find that the DM is indifferent between a certain 10 survivors on the one hand, and probability .2 of no losses and .8 of all lost. This tells us that $U[10 \text{ surv}] = 20$. Finding the indifference points for $i = 1, 2, 3$, etc. we can eventually plot all these utilities on a curve similar to the one shown in the next problem. With this utility function the commander can make rational, consistent choices between alternative courses of action.

PROBLEM: MOE's - Utility Function

You have a position to hold for which there are 12 possible approaches to be guarded. One man is adequate to guard one approach to your position. You have 20 men to start with. You have determined that your utility function, giving utility of various numbers of survivors from alternative defense plans, is as shown in the following chart⁸:

⁸This problem continues the third example in the text. Note that the simplest assumption might imply that the (constant) slope is the same in the two segments of the utility function. On the other hand, Lanchester's square law might be relevant here, leading to a quadratic (i.e. parabolic) shape in each segment. The actual curve shown here merely reflects the DM's judgement.

Utility Function



a) Calculate the utility of each of the following two strategies (show all calculations):

#1 -----p=.3, 3 survivors
 --- -----p=.6 12 survivors
 -----p=.1, 19 survivors

#2 -----p=.2, 0 survivors
 --- -----p=.8, 14 survivors

b) Which course of action will you select?

c) What ranking will an expected value D.M. give to these two strategies?

d) What ranking will a D.M. who uses $P[\text{surv.} \geq 12]$ give the strategies?

e) Classify the D.M. with the above utility function (the graph) as either risk averse, risk preferring, or indifferent, wrt this decision.

PROBLEM: The EW officer of a cruiser wishes to allocate his M disposable jammers to attackers. He doesn't know how many attackers will be sent against him, but he knows they'll be sent one at a time, one per day. He feels he has two, somewhat contradictory, goals. He's been told by the Captain that they must maintain their station as long as possible, but the Captain has also said that it's very important that they hold out for at least two more days. Let n be the number of full days the ship survives, i.e. the ship's position is abandoned on the $n+1$ -st day.

- a) What is the appropriate MOE? Is it $\Pr[n \geq 2]$, or $E(n)$, or $\Pr[n \geq 2] * E(n)$ or what? (be specific, with a rough sketch if possible)
- b) If the EW officer feels that it is 10 times as important to survive 2 days as to survive just one day, then sketch the MOE as a function of n [$0 \leq n \leq 10$].
- c) The EW officer has two possible jammer deployment strategies:
 - with strategy A : $\Pr[n=1] = .2$ and $\Pr[n=4] = .8$,
 - with strategy B : $\Pr[n=1] = .3$ and $\Pr[n=8] = .7$

which strategy will the EW officer prefer? Is he risk averse, preferring, or indifferent?

COST - EFFECTIVENESS

["The injury we do and the one we suffer are not weighed on the same scales.", Aesop - Fables]

Certainly, obtaining the most cost effective solution is an admirable goal, whatever that may mean. One sometimes hears people speak loosely of wanting to achieve the maximum benefit at the minimum cost. But, as a way of trying to factor costs into the MOE this is an inconsistent goal. One can attempt to:

1. maximize the benefit (e.g. $E[U(.)]$, P_K , $E[\text{no. enemy casualties}]$, etc.)
s.t. (subject to) costs $\leq$ budget

or

2. minimize the cost (e.g. total system cost or next incremental expenditure, etc.)
s.t. benefit $\geq$ requirements

or

3. maximize the benefit/cost ratio (i.e. "bang per buck" or spending efficiency)
s.t. cost $\leq$ budget & benefit $\geq$ requirements

or

4. maximize the "profit" (Benefit - Cost, if both can be expressed in dollars,
or some other common unit of measure)
s.t. cost $\leq$ budget

Each criterion could conceivably give a different rank ordering of alternatives., e.g. in a weapons system procurement. As with other incommensurable factors, one seeks to combine costs with lower level MOE's in a way that reflects the actual goals of the operation. This can only be decided case by case.

SONOBUOY EXAMPLE - REVISITED

The example on page 5 proposed a number of possible MOE's which were DE, at least when the outcomes were deterministic. We now show that even that is no longer true when benefits are divided by costs. Let $R_1 = 2\text{nm}$, and its cost, $C_1 = \$1000$.; $R_2 = 5\text{nm}$ and its cost, $C_2 = \$5000$. Clearly, the first sonobuoy has the greater value of benefit/cost ratio if "benefit" is defined as the sweep width, $2R$, i.e. $2R/C$ favors #1. (Presumably sonobuoy #1 has adequate performance, we can just buy more units to construct our barrier at a lower overall cost.) On the other hand, if benefit is defined as sweep area, πR^2 , then the second sonobuoy will have a 5 to 4 advantage over the first. The second sonobuoy is favored by evaluating $\pi R^2/C$.⁹ What this means is that you have to be just as careful in stating the real objective of the decision and identifying the corresponding MOE when dealing with Benefit/Cost ratios as when there are probabilistic outcomes.

⁹What has happened here is that in the previous, deterministic, analysis the MOE scale was arbitrary up to the monotonic functional transformation $h(\cdot)$, i.e. it was rubber, you could arbitrarily stretch some portions of the scale and shrink others without changing "greater than" into "less than". Once you divide by costs, however, you in effect define an absolute unit of scale (e.g. sweep width/\$ or area/\$). The issue becomes not just 'which has greater benefit', but 'exactly how much greater is it?' Hence we must select one, most appropriate, MOE for our benefit.

Earlier, it was only when we first considered sonobuoy range (R) to be a random variable that we were forced to consider quantities such as $p \cdot R + (1-p) \cdot R'$ or $p \cdot \pi R^2 + (1-p) \cdot \pi (R')^2$. This meant that we had to quantitatively compare the magnitudes of benefits on different parts of our MOE scale. This also forced us to select one, most appropriate, MOE.

The following example shows a number of alternative ways of solving (what is apparently) the same problem. How one decides in reality will depend on the situation and one's individual judgement.

EXAMPLE #5:

Given two possible new EW systems, A & B (anti-jamming, threat warning, etc.) which can be selected for installation on an existing aircraft (that would cost \$5M to replace) which one is the most "cost - effective"? The $P[\text{aircraft survival/mission}]$ with the present system is only .8 . The parameters for the new systems are:

	<u>System A</u>	<u>System B</u>
unit cost	\$100,000	\$500,000
$P[\text{survive/mission}]$	.9	.95

First analysis:

Rank the systems with respect to "no. missions/dollar";

the probability that the aircraft will survive n missions is given by a geometric distribution; $P(n) = p \cdot (1-p)^{n-1}$ where p is $1-.9$ or $1-.95$, the probability of being shot down per mission. Hence, the mean number of missions until the plane is shot down is $1/p = 10$ or 20 . Let's say we want to count complete missions, $E[\text{complete missions}] = 9$ or 19 (i.e. the plane is shot down on its last mission, possibly before it achieves its objective). Then for system A we might compute (WRONG!)

$$9/\$100K = 9\text{missions}/\$100K$$

and for system B we obtain

$$19/\$500K = 3.8\text{missions}/\$100K$$

which looks more favorable for system A. This may seem reasonable, since the \$5M cost of the plane is a "sunk cost", i.e. it's already been spent and we're just trying to achieve the maximum return for each additional dollar we spend.

On the other hand, we might take the point of view (why?) that the entire investment in airplane plus EW system should be included in the costs (and possibly also some imputed value for the crew, e.g. training costs). In that case we have for system A

$$9/\$5.1M = 1.77\text{missions}/\$1M$$

and for system B we obtain

$$19/\$5.5M = 3.45\text{missions}/\$1M.$$

Now system B looks more favorable, i.e. the ranking has reversed.

Second analysis:

Rank the systems with respect to max. benefit s.t. \$500K cost/plane.

Clearly system B is preferred since it yields 19 complete missions, on the average. This analysis assumes that we gain virtually no utility by spending under \$500K per EW system.

Third analysis:

Rank the systems with respect to min cost s.t. achieving ≥ 9 missions.

In this case system A wins since both EW systems provide adequate performance and A has the min cost. Here we gain virtually no utility by obtaining more than 9 missions. Perhaps the war will be over in 9 days (1 mission/day) or some important objective will be achieved with 9 missions. (There are other possible constraints on performance [benefit], e.g. one could set a minimum acceptable $P[\text{no. missions} \geq 9]$.)

Fourth analysis:

Max number of missions with 5 planes and only \$500K to spend:

there are two ways to spend our \$500K:

a) one plane with system B and 4 retain old system.

We obtain 19 missions from one plane and $4 \cdot (1/2 \cdot 1) = 16$ from the others, yielding a total of 35 missions, on average, from the 5 planes.

or,

b) 5 planes with system A, yielding $5 \cdot 9 = 45$ missions if we spend our money this way.

Hence, it is better to spend our \$500K on system A.

It is left as an exercise for the reader to try to develop reasonable scenarios in which each of the above analyses may be valid. The probabilistic analysis (or, in general, any mathematical modeling of an operational system) is not ambiguous or in question here. Rather, it is the complete definition of the problem and what the objective of the decision may be (and the correct definition of MOE).

Acknowledgement

The author would like to express his gratitude for helpful suggestions, material, and encouragement from Profs. Washburn, Eagle, Parry, and Hughes.

Selected References

1. Morse, Philip M. and Kimball, George E., Methods of Operations Research, OEG Report No. 54, (1950)
2. Hamburg, Morris, Statistical Analysis for Decision Making, chapter 13, Harcourt, Brace, Jovanovich, New York, (1977)
3. Raiffa, Howard, Decision Analysis, Addison-Wesley, Reading, MA (1968)
4. Baumol, W. J. , Economic Theory and Operations Analysis, Prentice Hall (1977)
5. Naval Operations Analysis, Naval Institute Press, Annapolis, (1978)
6. Hughes, Wayne P., Jr. (Editor), Military Modeling, The Military Operations Research Society (1984)