

Risk Tolerance Parametrics and the Maximal Value Frontier: The Value of Information for Risk-Averse Decision Making With Exponential Utility

Ronald E Davis*

College of Business, Marketing and Decision Sciences Department, San Jose State University, California, USA

*Corresponding author: Dr. Ronald E Davis, College of Business, Marketing and Decision Sciences Department, San Jose State University, California, USA, Tel: +374 10 23-72-61; E-mail: Ronald.Davis@sjsu.edu

Received date: March 27, 2014, Accepted date: October 24, 2014, Published date: November 08, 2014

Copyright: © 2014 Davis RE. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

When risk tolerance is varied from small to large values in a parametric programming study using CEV as the maxim and, a family of maximum value solutions are obtained that can be compactly summarized in policy region tabulations (in the discrete case) or asset allocation tabulations (in the continuous case). This gives the decision analyst an awareness of the full spectrum of solutions, from the one appropriate for the ultra-risk-averse decision maker, to the one appropriate for the risk neutral decision maker. When the interval of uncertainty from the least square risk tolerance estimation procedure is compared to the policy region tabulation or asset allocation tabulation, one obtains the solutions that are most preferred by the decision maker whose risk tolerance has just been estimated. From these solutions, all optimal in a neighborhood of the estimated risk tolerance, a final solution can be selected with complete awareness of what nearby optimal solutions would be like. In addition, for decision analyses involving an information option, parametric analysis reveals that the risk-averse CEVSI evaluation of information can be substantially greater than the risk neutral EVSI evaluation, which may lead to acquisition of information in cases where the risk neutral decision maker would pass it by. Hence use of EVSI for buy decisions about information may result in serious underutilization of information, leading to much greater downside risk than would be the case if CEVSI were used for those very same information buy decisions.

Keywords: Risk tolerance; Least square risk tolerance estimation; Parametric programming; Exponential utility; Certainty equivalent values; Value profiles; Maximal value frontier; EVPI and CEVPI; EVSI and CEVSI; Policy region tabulation; Asset allocation tabulation

Introduction

The use of exponential utility for the risk-averse decision analyses has been the default utility function family for many years by all manufacturers of computer software for this purpose. Later in this paper, it is shown that the reason for this selection is because it is only in this case that the traditional EVPI and EVSI values of information concepts from risk neutral theory can be extended in an unambiguous way to the risk-averse case (yielding the corresponding CEVPI and CEVSI concepts defined below). The key to this situation is the “delta-property” that the exponential utility possesses, presented by Howard Raiffa [1] and proved by Ronald Howard [2]. In particular, one has that the Certainty-Equivalent Value (henceforth shortened to Cash-Equivalent Value or CEV) of a shifted payoff distribution $X+\Delta$ satisfies $CEV(X+\Delta)=CEV(X)+\Delta$ just as for Expected Monetary Values one has $EMV(X+\Delta)=EMV(X)+\Delta$. No other nonlinear utility function form has this property.

The key parameter in the exponential utility function is the risk tolerance of the decision maker, the reciprocal of the risk aversion coefficient, which can be measured in terms of the dollar loss that will be tolerated in certain well-known calibration gambles. In this paper we present a general risk tolerance parametric programming method (or Risk Tolerance Parametrics, RTP) that consists basically of resolving a decision problem for a range of risk tolerance values and tabulating the scenario results for a variety of analyses. First we

develop some basic results pertaining to the limiting values of the CEV function as risk tolerance goes to zero or plus infinity. Then we derive the closed form formulas for the certainty equivalent functions for frequently use probability distributions that are particularly useful in practice. These functions permit the plotting of Value Profiles for different distributions that occur frequently in practice, which in turn allows the definition of the Maximal Value Frontier (MVF) [3] for a given set of gambles. The Risk Tolerance Parametrics methodology (referred to as RTP) required generating the Maximal Value Frontier and associated optimal policy choices becomes an integral part of a new risk-averse decision analysis paradigm. The RTP-MFV paradigm is akin to “getting the lay of the land” before building a house, or scoping out the size and shape of a forest before cutting down any trees. It is also closely related to the idea of developing a “requisite decision model” as described by Lawrence D. Phillips, by noting the behavior of one’s model not only in the vicinity of the final decision, but across the entire range of risk tolerances, from utterly risk averse (MaxiMIN criterion) to totally risk neutral (EMV criterion). In the portfolio optimization realm it is the generalization of the Mean-Variance Efficient Frontier methodology needed to accommodate asymmetric return distributions, such as the beta and gamma models presented in Davis and Davis [4,5].

In addition, valuable insights about the value of information can be gleaned from an appreciation of the full spectrum of MVF solutions before zeroing in on the final selection. We show how to extend the neutral EVPI and EVSI value of information concepts based on expected monetary values to the risk averse CEVPI and CEVSI concepts which are based on CEV values. A experience with numerous textbook examples shows that CEVSI tends to rise up higher than EVSI, in some cases to many times as much as EVSI, and then drops

back towards EVSI gradually, approaching EVSI in the limit as risk tolerance increases towards infinity. The fact that CEVSI > EVSI for a significant range of risk tolerances is very important from a practical standpoint because it means that information opportunities that would routinely be rejected on an EMV basis should actually be accepted on a CEVSI basis.

The presentation in this paper will be in six parts: (1) General Properties; (2) RTP-MVF methodology; (3) Simple gamble comparison example; (4) From EVPI and EVSI to CEVPI and CEVSI; (5) Portfolio Optimization example; (6) Conclusions. Appendices are provided that show (A) proof of the limits theorem that bounds CEV between the minimum and mean values of the payoff distribution; (B) derivation of CE value functions for uniform, histogram and normal distributions; (C) a compilation of CE value functions for most commonly used probability distributions. An earlier paper [3] presents the generalization of the value of information concepts EVPI and EVSI to the risk-averse case, yielding CEVPI and CEVSI. In the present paper we show that CEVSI can be much larger than EVSI in some risk tolerance ranges (on the order of 5 to 30 times or more), hence making the information options much more attractive to the risk averse than to the risk neutral decision maker.

General Properties

Exponential Utility functions may be written in the form $1 - \text{Exp}(-x/\tau)$ where x is a monetary amount and the scale factor τ , also a monetary amount with the same units of measure, is referred to as the Risk Tolerance of the decision maker. (Some books use "R" for this parameter, but it is a measure of risk tolerance, not of risk, so we prefer the Greek tau instead.). The CEV for a given distribution X with respect to a given risk tolerance τ is the solution to the equation stating that the utility of the CEV is equal to the expected utility of the distribution, or for exponential utility

$$1 - e^{(-CEV/\tau)} = 1 - E(e^{-X/\tau})$$

Solving for CEV yields the general result

$$CEV(X, \tau) = -\tau \cdot \ln[E(e^{-X/\tau})]$$

In the special case where X is a finite discrete distribution, this becomes

$$CEV(X, \tau) = -\tau \cdot \ln\left[\sum_{i=1}^n p_i e^{-x_i/\tau}\right]$$

This is the only form that is needed for most elementary decision analyses. Other formulas apply, of course, for other distributions used in more advanced analyses. A great number of these are presented in Appendix C to this paper.

In addition to the delta-property mentioned above, there are two limit theorems for Exponential Utility that are key to shape of the CEV function when plotted for a given distribution as a function of the Risk Tolerance parameter. This function is called the Value Profile for the distribution, and this function will tend towards the minimum payoff for the distribution (or $-\infty$ if the distribution is unbounded below) as risk tolerance approaches zero, whereas the CEV value function will tend towards the EMV or mean value of the payoff distribution as the risk tolerance approaches infinity. In fact the CEV Value Profile is always concave (for positive risk tolerances) and monotonically increasing between the two limiting values, so we have the following

Theorem: $\text{MIN}(X) \leq \text{CEV}(X, \tau) \leq \text{EMV}(X)$ for all risk tolerance values, where

$$\lim_{\tau \downarrow 0} \text{CEV}(X, \tau) = \text{MIN}(X) \quad \text{and} \quad \lim_{\tau \rightarrow \infty} \text{CEV}(X, \tau) = \text{EMV}(X)$$

The proof of this theorem is given in Appendix A. This leads immediately to the following corollary concerning the buying price $\text{BP}(X, \tau)$ and selling price $\text{SP}(X, \tau)$ of a gamble X .

Corollary: $\text{MIN}(X) < \text{BP}(X, \tau) < \text{EMV}(X, \tau) < \text{SP}(X, \tau) < \text{MAX}(X)$ for all $0 < \tau < \infty$

Here $\text{BP}(X, \tau)$ and $\text{SP}(X, \tau)$ are the buying and selling prices of a gamble X given by the following formulas:

$$\text{BP}(X, \tau) = \text{CEV}(X, \tau) \quad \text{and} \quad \text{SP}(X, \tau) = -\text{CEV}(-X, \tau)$$

Moreover,

$$\lim_{\tau \downarrow 0} \text{BP}(X, \tau) = \text{MIN}(X) \quad \text{and} \quad \lim_{\tau \downarrow 0} \text{SP}(X, \tau) = \text{MAX}(X)$$

And

$$\lim_{\tau \rightarrow \infty} \text{BP}(X, \tau) = \lim_{\tau \rightarrow \infty} \text{SP}(X, \tau) = \text{EMV}(X)$$

Proof: The buying price of X is that value (b) such that $\text{CEV}(X-b) = 0$. From the delta property one has $\text{CEV}(X-b, \tau) = \text{CEV}(X, \tau) - b = 0$ so that $b = \text{CEV}(X, \tau)$. The selling price of X is that value (s) such that $\text{CEV}(s-X) = 0$. Again the delta property leads one to $\text{CEV}(s-X, \tau) = s + \text{CEV}(-X, \tau) = 0$ so that $s = -\text{CEV}(-X, \tau)$. From the preceding theorem we know that $\text{MIN}(X) < \text{CEV}(X, \tau) < \text{EMV}(X)$ so $\text{MIN}(X) < \text{BP}(X, \tau) < \text{EMV}(X)$ as well. Likewise, $\text{MIN}(-X) = -\text{MAX}(X) < \text{CEV}(-X, \tau) < \text{EMV}(-X) = -\text{EMV}(X)$ so $\text{EMV}(X) < -\text{CEV}(-X, \tau) < \text{MAX}(X)$ or $\text{EMV}(X) < \text{SP}(X, \tau) < \text{MAX}(X)$

The limit results follow directly from the corresponding CEV results in the preceding theorem QED.

This result shows that the common misconception that buying and selling price are equal for exponential utility is actually false. In fact equality for buying and selling price is true only for the risk neutral decision maker with infinite risk tolerance. For any finite risk tolerance $\text{BP}(X, \tau) < \text{SP}(X, \tau)$ and in the most common case in which $2 * \text{MIN}(X) < \text{MAX}(X)$ there will exist a risk tolerance such that $\text{SP}(X, \tau) = 2 * \text{BP}(X, \tau)$ corresponding to the well-known maxim that "a bird in the hand is worth two in the bush."

Risk Tolerance Parametrics and Maximal Value Frontier

There are some general analysis procedures associated with the risk-averse decision analysis paradigm proposed here that we feel should always be carried out. The step by step is as follows:

- Using the Data Table command in Excel (or other equivalent programming tool) resolve the decision analysis problem over a grid of risk tolerance values ranging from suitably small to suitably large, storing the solutions obtained in a Scenario Results Table;
- Using the Goal Seek tool in Excel (or other equivalent programming tool) find accurate risk tolerance values where the solution changes (in the discrete case) or a solution breakpoint occurs (in the continuous case);
- Compile a Policy Region Tabulation to summarize the sequence of optimal policies obtained giving the applicable risk tolerance range for each (in the discrete case), or an Asset Allocation Breakpoint Table (in the continuous case).

- Plot the maximal CEV value obtained as a function of risk tolerance (defined as the MAXIMAL VALUE FRONTIER), with vertical lines locating policy changes or solution breakpoints, and annotation in each interval indicating what the optimal policy is in that interval;

- Perform a risk tolerance estimation Q&A with the decision maker to obtain a range of risk tolerance values containing the least square estimate of the risk tolerance for the decision maker in question;

- Superimpose the “interval of uncertainty” in the least square risk tolerance estimate on the policy region or breakpoint table or Maximal Value Frontier plot to identify those midrange policy or allocation solutions that obtain in the vicinity of the least square risk tolerance estimate;

- Present those midrange policies or allocations that occur in the “interval of uncertainty” to the decision maker for review and consideration.

- If model or parameter changes are obtained from the decision maker as a result of step 7, repeat those of the preceding steps that are necessary to arrive at the new results, and repeat step 7 again, until the decision maker is able to make a final policy or allocation selection.

The basic idea is to resolve the decision problem repeatedly on a grid of Risk Tolerance values (this is Risk Tolerance Parametrics or RTP) and from the resulting scenario results table, create a chart and table that shows the different optimal policies that result. In our view, this can and should be done before estimating the decision maker’s risk tolerance, to get a global view of the space of maximal value solutions before narrowing the range according to the risk tolerance estimation procedure. This is analogous to the practice in portfolio optimization, where the entire Mean-Variance Efficient Frontier is determined prior to selection of the optimal point on the curve by reference to some measure of the investor’s risk aversion, or as we prefer to do here, in terms of the investor’s risk tolerance. In fact, the RTP-MVF methodology described here can be thought of as a generalization of the well-known Mean-Variance Efficient Frontier concept (used for normally distributed returns) that can be used for any non-normal and asymmetric probability distributions. See Davis [5] for a treatment of portfolio optimization with asymmetric gamma distributed returns. A conference presentation [4] was also given on a portfolio model using asymmetric beta distributed returns.

Simple gamble comparison example

The RTP-MVF analysis paradigm just described can be illustrated with a simple gamble comparison example that also serves to acquaint the reader with the three CEV functions most often used in practice, or most likely to be used, as the methodology becomes more well-known. The first gamble has payoffs described by a finite discrete distribution for which the CEV formula has already been given. This will be option A for the example decision analysis, and will be a gamble having probabilities [0.3, 0.4, 0.3] for values [30, 60, 90] respectively. Hence the mean value is \$60, the variance is 540\$2 and the minimum payoff is \$30.

The second option B will be based on a histogram distribution, such as simulation and data analysis tools usually create. Such a distribution is described by n intervals and a set of n probabilities summing to 1.0 where interval i is the interval from xi-1 to xi and the interval endpoints [x0, x1, ...,xn] are arranged in increasing order. The probability for an observation in interval i is the given pi value, and the

distribution is uniform across each interval. For reference we note that the mean and variance for such a distribution are given by the following formulas, where the subscript H stands for the Histogram distribution.

$$\text{Histogram Mean} \quad \mu_H = \sum_{i=1}^n p_i \frac{(x_{i-1} + x_i)}{2} \quad \text{and}$$

$$\text{Histogram Variance} \quad \sigma_H^2 = \left(\sum_{i=1}^n p_i \frac{(x_{i-1}^2 + x_{i-1}x_i + x_i^2)}{3} \right) - \mu_H^2$$

The mean value is the weighted average of the interval midpoints, and the second moment is the weighted average of the interval second moments. The variance is given as the second moment less the square of the mean, as usual.

In Appendix C we show that the CEV for the continuous Histogram distribution takes the following form, where we use mi for the interval midpoints and wi for the interval half-widths.

$$m_i = (x_{i-1} + x_i) / 2 \quad ; \quad w_i = (x_i - x_{i-1}) / 2$$

$$CEV_H = -\tau * \ln \left(\sum_{i=1}^n p_i e^{-m_i / \tau} \frac{\sinh(w_i / \tau)}{(w_i / \tau)} \right)$$

This form is very similar to that used for the finite discrete distribution, where the interval midpoint takes the place of xi and there is an “adjustment factor” for width of the interval that involves the hyperbolic sine function, expressed in terms of relative half-width (i.e. relative to the risk tolerance). The adjustment factor approaches 1 as wi goes to zero. The data for option B is shown in the following Table 1.

Histogram parameter set

Breakpoints	Cum Probs
0	0
30	0.1
60	0.3
90	0.6
120	1

Table 1: Histogram Parameter Set.

This histogram has four intervals with interval probabilities [0.1, 0.2, 0.3, 0.4] respectively. The minimum value is zero, the mean value is \$75 and the variance is 975\$2. Compared to option A, its mean value is higher, but its variance is larger and the worst case is worse.

Finally the third option C is a gamble with a simple normal distribution with mean 80 and standard deviation 30. Hence the worst case is -∞ and the variance is 900\$2. It has the highest mean but the worst minimum value. The fairly well known CEV function for the normal distributions [6-10] is as follows.

$$CEV_N = \mu - 0.5\sigma^2 / \tau$$

The decision problem is to choose the option that has the greatest CEV. Obviously, the answer to the question depends on the size of the risk tolerance parameter. The RTP procedure requires we re-evaluate

the three options for a range of risk tolerance values. At the extremes, we find that based on expected value, the preference ordering would be $C > B > A$ whereas based on worst case, the preference ordering would be just the opposite, $A > B > C$. The risk tolerance parametrics (RTP) process entails resolving or reevaluating the problem on a grid of risk tolerance values which is such that the first optimal choice is A and the last optimal choice is C. This is easily accomplished in Excel using the Data/Table command. Plotting these three Value Profiles gives rise to the Maximal Value Frontier and a set of policy regions (intervals in risk tolerance) in which the optimal choice remains constant as shown in the Figure 1. As expected, we find that there is no one optimal solution for all risk tolerance values, rather there is a set of options which are optimal for some risk tolerance intervals, and also a set of options which are not optimal for any risk tolerance value. Here is the chart depicting this analysis for the present case.




Figure 1: Value Profile Comparison.

The profile for option A begins at 30 (the min for A) on the left and rises towards a limiting value of 60 (the EMV for A) on the right. The profile for option B begins at zero and rises towards a limiting value of 75 on the right. The profile for option C begins at $-\infty$ and rises towards a limiting value of 80 on the right. The upper envelope of the collection of value profiles is the Maximal Value Frontier for the problem. In this case, A is optimal for all risk tolerances greater than zero and less than 12.28 (approx.). For risk tolerances greater than the policy region breakpoint value 12.28 it is seen that option C is preferred. In this case, option B is never optimal since it does not appear in the maximal value frontier, meaning that CEV (B) is always less than the maximum of CEV (A) and CEV(C). Hence we can define a new sort of dominance, dominance in CE-Value. We say that B is dominated by {A, C} in CE-Value when $CEV(B, \tau) < \text{MAX}(CEV(A, \tau), CEV(C, \tau))$ for all positive risk tolerances.

The remaining element of the RTP-MVF procedure to be shown is the Policy Region Tabulation (PRT). In this table the risk tolerance intervals are listed in increasing order together with the “policy” that is optimal on each risk tolerance interval. The word policy is used because this methodology can be applied to large decision trees having many decision nodes, and a new interval occurs whenever any one of the decisions in the decision rule changes. In our simple example, there is only one decision node, so the policies are specified by the alternative chosen. In a more complex situation, a more complex decision rule would be associated with each region in the tabulation.

These decision rules would be computed by using CEV at each chance node instead of the usual risk neutral EMV evaluation.

By subtracting CEV (A) from CEV(C) we can search for the zero point for the difference (using the Goal Seeker Tool in Excel, for example). In this case the (approx) risk tolerance obtained is 12.28 (dollars, or other monetary unit as defined in the problem). Hence the PRT in this case is simply shown in the Table 2.

0 to 12.28	Choose A
12.28 to $+\infty$	Choose C

Table 2: Policy Region Tabulation.

One finds that if option B were to be enhanced appropriately, it enters into the PRT and the MVF as the optimal choice in some midrange risk tolerance interval. In fact, if the two lower breakpoints are moved up to 15 and 45 respectively (leaving the other parameters unchanged), then the mean of B moves up to \$78, variance drops to 718.5\$2 and the Maximum Value Frontier appears as in the chart below shown in the Figure 2.




Figure 2: Revised Value Profile Comparison.

The associated Policy Region Table 3 in this case becomes

0 to 5.31	Choose A
5.32 to 40.17	Choose B
40.18 to $+\infty$	Choose C

Table 3: Policy Region Table.

Now there are no dominated alternatives; indeed, all options are optimal over some range of risk tolerances. A, being the MaxiMIN choice, is still optimal for the smallest risk tolerance range, and C, being the EMV choice, is still optimal for large risk tolerance range. But now there is a significant midrange for risk tolerance in which B is the preferred choice. For larger more complex problems, the intermediate solutions between the MaxiMIN choice and the EMV choice may be more numerous, of course. The first major purpose of the RTP-MVF methodology is to find out what these intermediate solutions are, their associated risk tolerance ranges, and lay them “on the table” for explicit consideration. Then estimation of the appropriate risk tolerance range will identify which solutions on the

MVF need to be considered most closely to make a final selection. Risk Tolerance estimation is covered in later in section 6.

Decision Tree Problem with Information Option

The standard value of information concepts EVPI & EVSI for risk neutral evaluations have natural extensions in the risk-averse case, denoted CEVPI and CEVSI. Instead of measuring an increase in Expected Monetary Value ($EMV = \sum p_i x_i$), one measures an increase in Cash Equivalent Value ($CEV = -\tau \ln(\sum p_i \exp(-x_i/\tau))$). This is simple enough to define, but some surprising and significant things happen when you carry out the risk tolerance parametric analysis in specific cases. The example developed in this paper shows, in particular, that CEVSI can be MUCH LARGER than EVSI in certain mid-range risk tolerance intervals, meaning that the decision to buy or not buy information may be different as well. In the example below, the CEVSI/EVSI ratio rises to over 5.5, so that in many cases the information option that might be passed over by the risk neutral decision maker should in fact be purchased by the risk-averse decision maker. Other textbook examples have been seen in which the value ratio is over 27.

There is another extremely important point that can be illustrated with this same example. This point constitutes one of the principal justifications for accepting the delta-property axiom and therefore using exponential utility for risk-averse analyses. It is the consistency between the backwards induction process and the value of the information. Since the value of the information is computed as if it were free, the cost of the information does not enter into either EVSI or CEVSI. Let us suppose that the CEVSI obtained for a given risk tolerance exceeds the cost of the information, indicating that the information should be bought. If the cost of information is now deducted from all terminal payoffs on branches following the decision to buy, the backwards induction can be done again. We would like to see that the optimal policy (and value) obtained taking the cost of information into account agrees with the optimal policy (and value) indicated by the CEVSI computation. In fact, if we do not get the same policy (and same value), then we seem to have a contradiction that is hard to explain. It turns out that the only way to avoid this kind of contradiction is to require that the delta-property continue to hold in the risk-averse case as it does for the EMV case. And as we have seen, this means that the utility functions for risk-averse analysis must come from the exponential utility family.

The example analysis

ACE Computer Company has been using the Be-Sure Survey Company to predict the success of new products. Over a period of years ACE has found that when a new product was successful, i.e., sales were high for that product, the survey Co. study had predicted success 30% of the time, showed inconclusive results 60% of the time and predicted failure 10% of the time. The record also indicated that when sales for a new product were low Be-Sure Survey Co. predicted success 10% of the time, showed inconclusive results 40% of the time and predicted failure 50% of the time. ACE Company has established the probability of high sales for a new product at 40% and low sales have a 60% probability.

It will cost ACE Company \$1 million to introduce its new product, and if Be-Sure Survey is retained again, it will cost \$100 thousand for the survey. If sales are high they expect to gross \$4 million and they would expect to gross \$0.5 million on low sales.

Prior analysis

We first construct a payoff table for the “main” decision, which is whether or not to market the product. If the product is not marketed, there is no introduction cost and no revenue, so the payoff is zero regardless of the level of potential demand. If the product is marketed, the net profit is 4 - 1 or \$3 million under the High demand scenario and 0.5 - 1 or a loss of (\$0.5) million under the Low demand scenario as shown in the Table 4.

	A0	A1	V*
E1 (High demand) 0.4	0	3	3
E2 (Low demand) 0.6	0	-0.5	0
EMV	0	0.9	1.2

Table 4: Prior Analysis Go/Nogo Payoff Table.

The expected value is higher for the A1 decision (market the product). Hence the highest expected value that can be achieved without perfect or sample information is \$0.9 million or \$900,000. With perfect information available, the expected value increases to \$1.2 million or \$1,200,000. EVPI is the difference, or \$300,000. Since Be-Sure Survey Company is only asking \$100,000 for their survey, there is some possibility, at least, that it might be worthwhile to use their services. But the decision about this issue must await the outcome of the decision tree analyses of sample information described next.

Posterior analysis

Since the EVPI for this situation exceeds the asking price of \$100,000 for the Be-Sure Survey marketing study, it is conceivable that the market study might be worthwhile. But this depends upon the track record, which Be-Sure has had in previous studies of the same sort in the past. From the data given in the problem statement we can tabulate the following prior and conditional survey result probabilities shown in the Table 5:

Prior Probability	State	P(PS Ei)	P(I Ei)	P(PF Ei)	SUM
0.4	High Demand	0.5	0.4	0.1	1
0.6	Low Demand	0.1	0.4	0.5	1

Table 5: Conditional Survey Result Probability Table.

State	P(PS and Ei)	P(I and Ei)	P(PF and Ei)	Marginal
High Demand	0.2	0.16	0.04	P(E1)=.4
Low Demand	0.06	0.24	0.3	P(E2)=.6
Marginal	P(PS)=.26	P(I)=.40	P(PF)=.34	1

Table 6: Joint Probability Table.

The row sums are 1.0, indicating that each row is a separate probability distribution, conditioned on which demand level applies. The column sums need not be 1.0. Now, to form the joint probability table, we must multiply each conditional probability by the prior probability at the beginning of the row in which it occurs. This yields the following shown in the Table 6:

Now the row sums give the marginal probabilities for the demand levels, i.e. the prior probabilities in this case, and the column sums give the marginal probabilities for the survey results. Finally, by dividing the joint probabilities in each column by the marginal survey result probability at the base of the column, we get the “posterior” probabilities, or the conditional probabilities for the demand levels given the survey result shown in the Table 7.

State	P(Ei PS)	P(Ei I)	P(Ei PF)
High Demand	.2/.26=10/13	.16/.40=0.4	.04/.34=2/17
Low Demand	.06/.26=3/13	.24/.40=0.6	.30/.34=15/17
SUM	1	1	1

Table 7: Posterior Probability Table.

In this table the columns sum to one, and the row sums need not be one. Each column gives the “revised” or “updated” or “posterior” probability distribution for the demand level, given the survey result reported by Be-Sure Survey Co. Notice also that if Be-Sure predicts success (PS) then the probability of High Demand increases from its

prior of 0.4 to a posterior of almost 77%. On the other hand, if Be-Sure predicts failure (PF) then the probability of High Demand decreases from its prior of 0.4 to a posterior of less than 12%. And if Be-Sure is indeterminate (I) about sales, then the decision-maker regards High Demand as a 40-60 proposition, just the same as the prior probabilities.

Pre-posterior analysis

Our next task is to evaluate the EVSI (Expected Value of Sample Information) for the Be-Sure survey result. We wish to know what it would be worth to us in increased expected value, if the information were free. By comparing this EVSI with the price Be-Sure is asking for its prediction (\$100,000) we can determine whether or not to buy the information. In particular, we can compute the ENGSI (Expected Net Gain of Sample Information) which is the difference between the two, $ENGSI = EVSI - \text{Cost of Information}$. If there were more than one Survey Company we were considering, each with a different track record and a different cost of survey, we could compute the ENGSI for each alternative information source. In this case we could pick that one which gives the largest expected net gain, if any is positive, or make the main decision without surveying if they are all negative.

In order to evaluate EVSI, we need to develop and “roll-back” the decision tree corresponding to the “BUY SURVEY” decision. In this analysis, the cost of the survey will be neglected (i.e. treated as zero) and the marginal survey result probabilities and the posterior demand level probabilities will be employed, as shown in the tree below as Figure 3.




Figure 3: Result probabilities and the posterior demand level probabilities.

We have computed the expected revenues at the end of the tree, and then subtracted the \$1 million cost of marketing only when the Market decision yields the higher net return. The net return figures are shown in each decision box based on the optimal policy from that point forward. In this case the optimal decision policy is to market the product if the Be-Sure result is PS or I, and not market the product if the Be-Sure result is PF.

The net expected values for each survey result are then weighted by the survey result probabilities. This gives an expected value of \$0.93 million given (free) sample information, or \$930,000. When this is compared with the best we could do using prior information, namely \$900,000, we have

$$EVSI = EV|SI - \text{Max EMV}(A_i) = \$930,000 - \$900,000 = \$30,000$$

The expected value of having access to the Be-Sure Survey Co. result prior to the marketing decision is only \$30,000 and they are asking \$100,000 for it. Thus based on expected values, the ENGSI comes out to be a large negative amount (\$70,000). If the decision-maker is risk neutral (i.e. makes decisions based on expected values only), then the most that could be paid for the Be-Sure survey result would be \$30,000. Paying anything more would cause the expected net gain to go negative, and hence is inferior to simply marketing the product without the benefit of the survey result.

The Risk-Averse Analysis Using a Risk Tolerance

In reality, we know that decision-makers generally have a certain amount of risk aversion, so that they exhibit some sensitivity to the worst case result as well as the expected value of the outcome. The certainty equivalent value (or CEV, for short) of a chance outcome is therefore obtained by discounting the expected value of the outcome by a certain "risk premium" which depends upon the risk tolerance of the decision-maker. The optimal decision rule and also the value of perfect and survey information must therefore be computed in terms of the two cash equivalent values involved, assuming the information is free, so that we have the following Table 8.

$CEVPI_{\tau} = CEV_{\tau} PI - \text{Max}\{CE_{\tau}(A_i)\}$
$CEVSI_{\tau} = CEV_{\tau} SI - \text{Max}\{CE_{\tau}(A_i)\}$

Table 8: Information Value Concept Definitions.

As the risk tolerance scale factor τ decreases from ∞ towards 0, the policies, both with and without sample information, may change to reflect a progressively more and more risk-averse posture, until the risk tolerance is so small that the product would not be marketed under any conditions. At this point the CEVSI drops to 0. Our problem in this risk tolerance parametric analysis is to determine at exactly which values for risk tolerance does the policy change (with or without survey information). Also, how do $CEV_{\tau}|PI$, $CEV_{\tau}|SI$ and $\text{Max}\{CE_{\tau}(A_i)\}$ and thus $CEVPI_{\tau}$, and $CEVSI_{\tau}$ change between these "breakpoints" in the risk tolerance level. When we plot these cash equivalent values between the policy change "breakpoints" in different colors, we get what is called a "Rainbow Diagram." EXCEL spreadsheets for accomplishing this, and the associated charts of CEVPI and CEVSI, are shown below, but first we illustrate the computational process using just one particular value for τ , namely $\tau = 1$ million \$ (i.e. \$1,000,000).

Prior analysis with a risk tolerance

Let's calculate the CE value for the prior analysis first, without using the results of the survey. We have

$$CE(A1) = -1 \cdot \ln[0.4 \text{EXP}(-4/1) + 0.6 \text{EXP}(-0.5/1)] - 1 = -0.009106 = -\$9,106$$

where the prior probabilities have been used to evaluate the chance node on level of demand. We would get the same value if we had evaluated the net profit figures from the payoff table,

$$CE(A1) = -1 \cdot \ln[0.4 \text{EXP}(-3/1) + 0.6 \text{EXP}(0.5/1)] = -0.009106 = -\$9,106$$

This is due to the "Value Additivity" or delta-property of Exponential Utility mentioned before. We can compute the CE value ignoring the 1M cost and then subtract the 1M, or we can subtract the

1M cost first and then compute the CE value; we get the same answer either way. Also note that the value of the A1 alternative has dropped from its former \$900,000 level all the way down to just under zero, so the former GO with the product is replaced with a NOGO preference. The Risk Premium in this case is given by

$$RP = EMV - CE_{\tau} = \$900,000 - (-\$9,106) = \$909,106$$

Notice also that the value of the V^* gamble (given Perfect Information) has changed as well, since we have

$$CE_{\tau}|PI = -1 \cdot \ln[0.4 \text{EXP}(-3/1) + 0.6 \text{EXP}(0/1)] = .478173185 = \$478,173.185$$

Consequently, in this case the value of perfect information is given by

$$CEVPI_{\tau} = CEV_{\tau}|PI - \text{Max}\{CE_{\tau}(A_i)\} = \$478,173.19 - \$0 = \$478,173.19$$

This is substantially greater than the EVPI obtained before, namely, \$300,000. Thus the value of perfect information can be worth more to the risk-averse decision-maker than to the risk neutral decision-maker. We shall see shortly that the same is true for the value of sample information. Also by varying the risk tolerance through a range of values, we get a chart of CEVPI as a function of τ .

Posterior analysis with a risk tolerance

Let's turn now to the computation of the value of sample information. When the backwards induction process is carried out for a risk-averse decision maker, ALL chance nodes are regarded as gambles at which a cash equivalent value must be computed from the probabilities and the cash equivalent values for outcomes at a the node. Thus all expected value computations are replaced with cash equivalent value evaluations. This must be done THROUGHOUT the ENTIRE TREE, not just at the ends of the tree. Let's see how this works out for the Be-Sure decision with $\tau = 1$.

Consider the three demand level chance nodes using the three sets of posterior probabilities. After the "predict success" result (PS), the posterior probabilities are 10/13 and 3/13, which lead to a net value of

$$CE|PS = -1 \cdot \ln[(10/13) \text{EXP}(-3/1) + (3/13) \text{EXP}(0.5/1)] = 0.870428936 = \$870,428.94$$

Likewise, after the indeterminate result (I) we have

$$CE|I = -1 \cdot \ln[0.4 \text{EXP}(-3/1) + 0.6 \text{EXP}(0.5/1)] = -0.009106 = -\$9,106 < 0$$

as previously obtained in the Prior Analysis, so the decision is NOGO in this case with a net value of 0. And after the "predict failure" result (PF) we have

$$CE|PF = -1 \cdot \ln[(2/17) \text{EXP}(-3/1) + (15/17) \text{EXP}(0.5/1)] = -0.37885509 < 0$$

so in the last case it's also preferable not to market the product, and the net value is 0.

Now backing up to the chance node for the survey result, we again use the cash equivalent formula (NOT an expected value calculation) and obtain

$$CE|SI = -1 \cdot \ln[0.26 \text{EXP}(-0.870428936/1) + 0.4 + 0.34] = 0.163836632 = \$163,836.63$$

where the last two exponential terms reduce to 1.0 since the payoff value in that case is zero and $\exp(0)=1$. Placing these cash equivalent values on the decision tree gives us the following modified Figure 4:

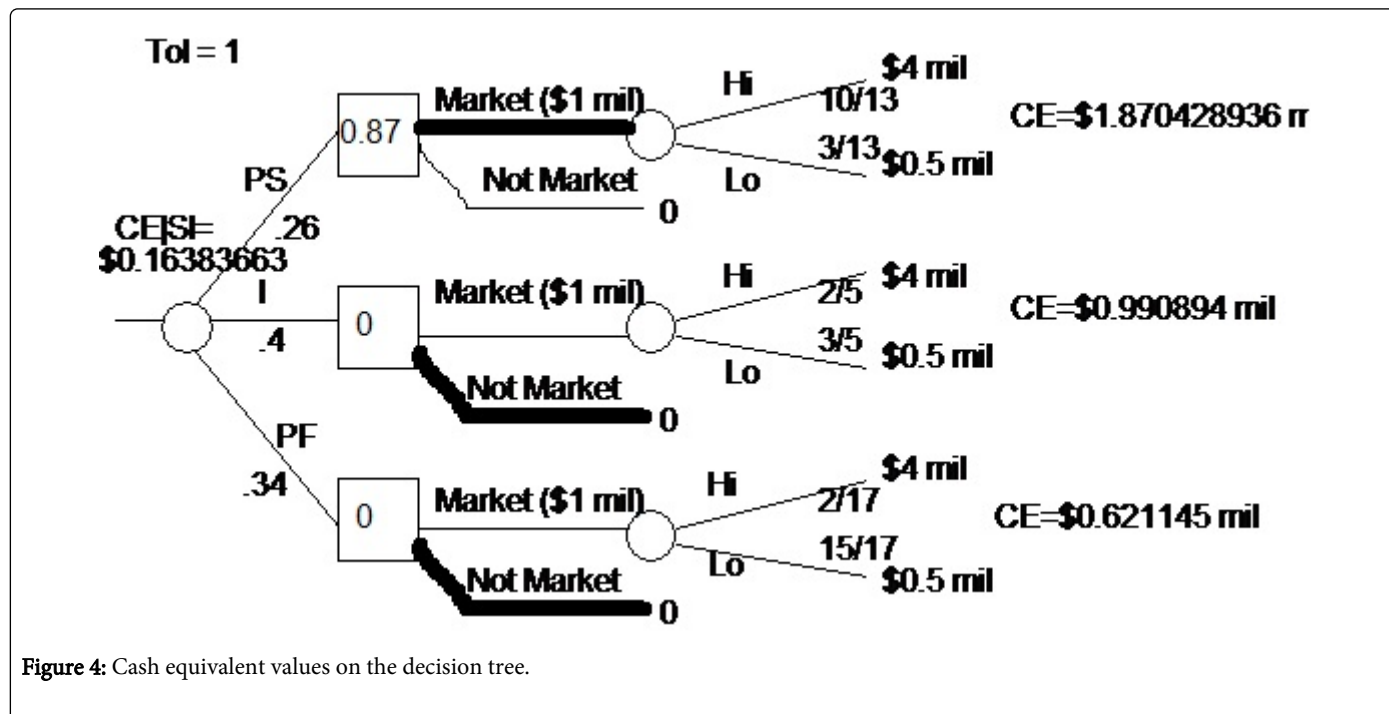


Figure 4: Cash equivalent values on the decision tree.

And the corresponding tree for the prior analysis has changed to:

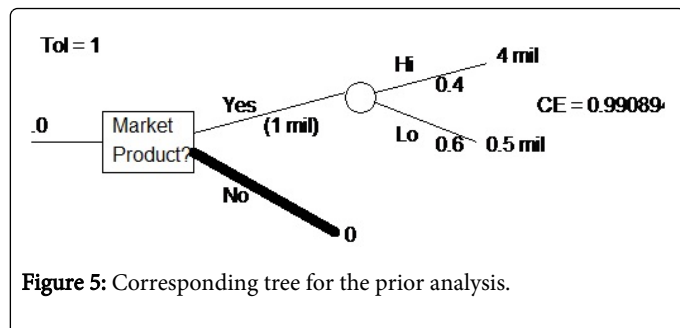


Figure 5: Corresponding tree for the prior analysis.

Finally taking the difference between $CEV|SI$ and $CEV|priors$ we obtain

$$CEVSI = CE|SI - \text{Max } CEV(A_i) = \$163,836.63 - \$0 = \$163,836.63.$$

Now this IS interesting! The CEVSI has gone UP dramatically from the EVSI, which was only \$30,000. The information is worth MUCH MORE to the risk-averse decision-maker than to the risk neutral one. Hence we cannot assume that the most one should pay for the sample information is \$30,000. We have shown that when $\tau = 1$, the decision-maker would be willing to pay up to \$163,836.63 for the survey result, a full \$130,000 more than the risk neutral decision-maker would. Hence the \$100,000 asking price is seen as attractive in this case, and $CENGSI = \$163,836.63 - \$100,000 = \$63,836.63$. Also, note the magnitude of the increase in information value in terms of the ratio $CEVSI/EVSI = 5.46$, nearly five and a half times as large. In the course of instruction, we have seen this ratio as high as 27 times higher.

Risk Tolerance Parametrics

By varying the risk tolerance across a range of values, one can easily show that in fact there is a substantial range of risk tolerances that would justify the \$100,000 price for the sample information. In fact, as you systematically vary τ , you will find that there are three distinct “breakpoints” in the analysis where either the prior or the posterior policy changes. As you increase τ from 0, you will first notice that the Max CEV for the prior policy is zero, which means that the product is not introduced, based on prior information only. At a somewhat larger value, one finds that the Be-Sure Survey option becomes attractive, and that an I result from Be-Sure is not sufficient, and one only goes ahead with the product if the Be-Sure result is PS. In the next policy region, a PS or an I result is sufficient to justify going ahead with the product. And finally, the CEVSI drops below \$100,000 again so the optimal decision is to just introduce the product with no survey, which is the risk neutral or EMV policy. From the attached graph of CEVSI, you can see that it varies nonlinearly in a smooth way between breakpoints, and that it achieves a maximum value at a unique value of τ . What is the value of τ which maximizes CEVSI? What is the maximum value that CEVSI attains? What is the maximum value that CEVSI/EVSI attains? Explain the significance of this result.

To summarize the results of the parametric analysis on risk tolerance, it is convenient to create a table showing the optimal policies for each policy range, and the interval of risk tolerances over which that policy is optimal. The policies that appear in this table constitute the “MAXIMAL VALUE FRONTIER” for the problem, and all other policies are said to be CEV-dominated by these which are optimal for one risk tolerance or another. A plot of the maximum CEV for each risk tolerance in which the area under the curve is color coded according to which policy is optimal is called a Rainbow Diagram for the Analysis and depicts the range of optimality for each policy

graphically. The table and chart for the Be-Sure Survey analysis are shown below. The color coding in the third Column correlates with the color coding on the Rainbow Figure 6.




Figure 6: Color Coded Maximal Value Frontier.

This is associated with the following

POLICY	Risk Tolerance Range	Color Code
Don't Survey, Don't Market	0 - \$694,576	Yellow
Survey, Market on PS only	\$694,577 - \$1,021,720	Green
Survey, Market on PS or I	\$1,021,721 - \$3,199,662	Cyan
Don't Survey, Just Market	\$3,199,663 - +INF	Purple

Table 9: Policy Region Tabulation.

The first policy is appropriate for the ultra-risk-averse decision maker who evaluates gambles in terms of their worst case shown in the Table 9. And the last policy is appropriate for the risk neutral decision maker who evaluates gambles in terms of their EMV. The benefit of the risk tolerance parametric analysis is that it shows up two more “in

between” policies that are optimal for mid-range risk tolerance levels, both of which indicate purchase of the Be-Sure Survey result. Note that the risk tolerance range in which the Survey is purchased extends from \$694,577 all the way up to \$3,199,662. This is a significant interval that might very well include the risk tolerance appropriate for the ABC Computer Company executives. Hence the information option cannot be rejected just because the \$100,000 cost exceeds the EVSI of \$30,000. If the decision makers are risk averse, as they usually are, then one must compute the CEVSI values, and these may well indicate purchase of the information even when EVSI does not.

To get a visual sense of the degree to which CEVSI may exceed EVSI, one can construct the plot of CEVSI versus risk tolerance, as shown in the Figure 7 below.

In this case the EVSI is only \$0.03 Million whereas the CEVSI rises to over \$0.16 Million when risk tolerance is in the neighborhood of \$1 Million, an increase of over five times. This makes the information attractive (even when priced at \$100,000) across a significant risk tolerance range, whereas the EVSI result would reject the survey option out of hand.

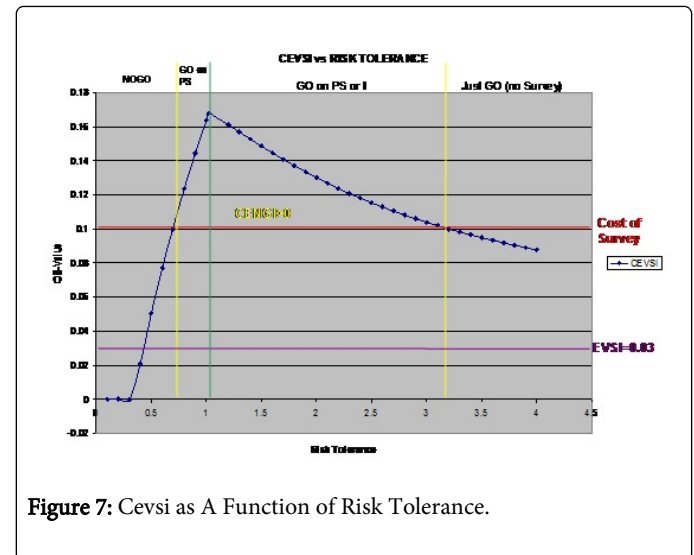


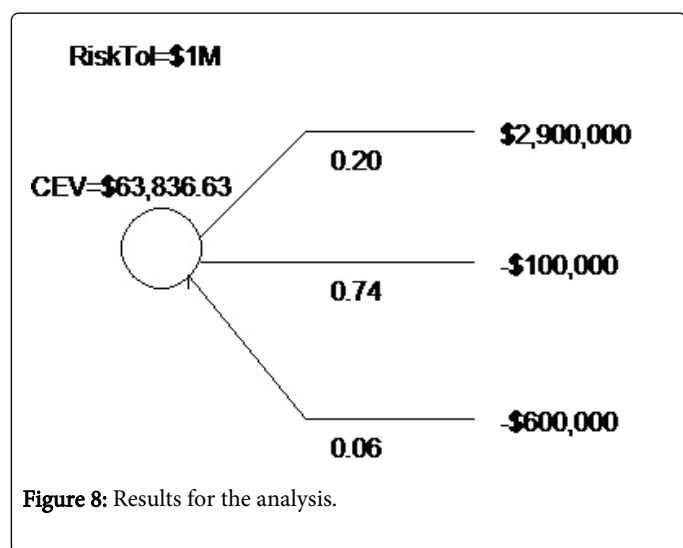
Figure 7: Cevsi as A Function of Risk Tolerance.

Consistency check

One of the nice features of risk neutral EMV analysis is that the EMV of the optimal policy risk profile is equal to the EMV developed by the backwards induction process which shows over the first node in the decision tree. This is insured because EMV satisfies the delta-property requirement that $EMV(X+\Delta) = EMV(X)+\Delta$ where Δ is any constant. It would be nice if this remained true for risk-averse analyses as well, and, in fact it does remain true if the utility function also satisfies the delta-property so that $CEV(X+\Delta) = CEV(X) + \Delta$. For other utility functions that do not have this property, contradictions may arise in which the value of information results are different from what is obtained via the backwards induction process. That is, the results obtained by keeping the cost of information on the branch preceding the sample information chance node may be different from the results obtained by netting out the cost of the sample from the payoffs at the end of the tree. If this occurs, then one is in a quandary to explain why the optimal solution from the backwards induction is not the one indicated by the value of information results. The only way to avoid this quandary is to require the utility function to have the

delta-property, which, as we know, implies that it be from the exponential utility family.

Let us now confirm equality of the results for the analysis completed in this example with risk tolerance at \$1 million as before. By deducting the \$100,000 survey cost from the CEVSI we obtained the value \$63,836.63 for the optimal policy, which was to “Buy Survey; Market only if PS”. Now we will develop the risk profile for this policy by collapsing the tree down to a single chance node, deducting the \$100,000 survey cost from the affected terminal node values. We find there is a 74% chance of just losing the \$100,000 cost of the sample, a 20% chance of the “big hit” being \$2.9 million in this case, and only a 6% chance of taking a major loss of \$600,000. Hence in Figure 8, we have



Note that the CEV of the risk profile for the optimal policy is EXACTLY equal to the CEV we got by backwards induction on the decision tree with the cost of information subtracted only once at the beginning of the tree, not multiple times at the end of the tree. The policy obtained is the same as well, because the choices made at each of the decision nodes will be the same in either analysis. This equality of results will always be true for exponential utility analysis because of the delta-property that is true for this family of utility functions.

Portfolio optimization example

The standard Markowitz portfolio optimization problem is usually formulated with a minimum variance objective and an expected return constraint. Parametrics are performed by increasing the expected return target from the minimum possible to the maximum possible to sweep out the mean-variance efficient frontier. In this context, we prefer to formulate the objective function as a maximization of the certainty equivalent value of the portfolio return distribution. The parametrics are then done by varying the risk tolerance parameter from 0 to infinity, obtaining a minimum variance solution at one end and a maximum expected value at the other. Specifically we

$$\text{Maximize } r'x - 0.5x'Cx/\tau \text{ subject to } 1'x=1 \text{ (or } 100) \text{ and } x \geq 0$$


Here r is the column vector of expected returns, x is the column vector of investment levels (normally fractions summing to 1, or percentages summing to 100, or dollar investments summing to the portfolio fund size). C is the covariance matrix for the returns, and $1'$ is a row vector of ones the same dimension as x .

It can be shown that when the parametric analysis on risk tolerance is done in this case, the asset allocations turn out to be piecewise linear, with breakpoints where an allocation declines to zero, or rises from zero. This is a particular convenient result since the entire range of solutions can then be obtained by linear interpolation in a breakpoint table showing the allocations only at breakpoints. Since the number of breakpoints is usually a fairly small number, this is an extremely compact and useful representation of all the solutions on the MVF.

In this example, based on data given in [5], the first stock is for APD, the second is for IBM, and the third is for XON. The returns and covariance matrix for the example are

$$r = \begin{bmatrix} 0.0988 \\ 0.1487 \\ 0.1756 \end{bmatrix} \quad C = \begin{bmatrix} 0.0264 & 0.0036 & 0.0336 \\ 0.0036 & 0.0351 & -.0029 \\ 0.0336 & -.0029 & 0.0558 \end{bmatrix}$$

The mean returns and variances increase in the sequence given. Essentially the same results were obtained using the Excel Solver and the MATLAB Optimization Toolbox quadprog function, as shown on the Figure 9 below.



For small risk tolerances (less than 1.1.6193%) XON remains at zero level, having the largest variance, with about equal allocations to APD and IBM. On this first interval, the allocation shifts from APD to IBM with increasing risk tolerance since the expected return for IBM is greater than for APD. On the next interval, XON comes into the solution and the shift to IBM increases until APD goes to zero at $\tau=6.9434\%$. On the third interval the allocation shifts from IBM to XON with increasing risk tolerance since the expected return of XON exceeds that of IBM. Eventually at $\tau=218.2156\%$ the IBM allocation goes to zero, and 100% of the allocation goes to XON for all larger risk tolerances. Since (it can be shown) the allocation plots are piecewise linear for each investment class, the entire solution can be characterized in a table listing the allocations at each breakpoint in the analysis, where an activity that had been zero becomes positive or where an activity that had been positive becomes zero. In this case that table is:

On the intervals between breakpoints, the allocation formulas are linear in risk tolerance shown in the Table 10; hence intermediate

solutions can be obtained with linear interpolation between successive breakpoint solutions. For this particular example the linear formulas are as follows:

Risk Tolerance %	APD %	IBM %	XON %
0	58.011	41.989	0
1.6193	56.5229	43.4771	0
6.9434	0	58.7717	41.2283
218.216	0	0	100

Table 10: Allocation Percentages Vs Risk Tolerance.

On [0, 1.6193]

$$\begin{bmatrix} APD \\ IBM \\ XON \end{bmatrix} = \begin{bmatrix} 58.011 \\ 41.989 \\ 0 \end{bmatrix} + \begin{bmatrix} -0.918969 \\ 0.918969 \\ 0 \end{bmatrix} \tau$$

On [1.6193, 6.9434]

$$\begin{bmatrix} APD \\ IBM \\ XON \end{bmatrix} = \begin{bmatrix} 73.7145 \\ 38.8252 \\ -12.5397 \end{bmatrix} + \begin{bmatrix} -10.6165 \\ 2.8727 \\ 7.7438 \end{bmatrix} \tau$$

On [6.9434, 218.2156]

$$\begin{bmatrix} APD \\ IBM \\ XON \end{bmatrix} = \begin{bmatrix} 0 \\ 60.7032 \\ 39.2968 \end{bmatrix} + \begin{bmatrix} 0 \\ -0.2782 \\ 0.2782 \end{bmatrix} \tau$$

These formulas are shown to emphasize the fact that asset allocations vary linearly with risk tolerance between breakpoints. In practice, linear interpolation in the former breakpoint table will suffice.

Another important advantage of the maximum CE value formulation is that it generalizes to asymmetric distributions for which the mean-variance efficient frontier is no long optimal in any maximum expected utility sense. To illustrate this point, let us now suppose that the three given gambles in the earlier simple example (A, B, and C) are the forecasted total dollar returns for a \$50 investment in each one (we use the enhanced version of Gamble B for this example). Let us also assume that this is a mutual fund type situation that allows for fractional allocations in each gamble. When it is possible to split the \$50 investment between the three gambles, then we obtain a portfolio optimization problem that is formulated in terms of maximizing the portfolio CEV which is a function of the fractions of the investments in each gamble. If the return distributions are independent one from the other, then, because of the value additivity property, we have that $CEV(\text{Portfolio}) = CEV(fa \cdot A) + CEV(fb \cdot B) + CEV(fc \cdot C)$ where the fractions (fa, fb, and fc) are nonnegative and sum to 1. Formally, the problem is to

$$\text{Maximize Portfolio CEV} = CEV(fa \cdot A) + CEV(fb \cdot B) + CEV(fc \cdot C)$$

$$\text{Subject to } fa + fb + fc = 1 \text{ and } fa, fb, fc \geq 0$$

Scaling Gambles A and B is accomplished by scaling the xi parameters in their definitions, leaving the probability assignments the same. For the normal distribution, the mean and standard deviation are scaled by the fractional allocation fc. Maximizing Portfolio CEV subject to the given constraints is a well formed and well behaved

nonlinear programming problem which is easily solved for a range of risk tolerance values, using the Solver Table Excel Add-in for example.

We can anticipate the behavior of the optimal allocation at the two extremes, i.e. where risk tolerance tends towards 0 and where it tends towards infinity. Since Gamble A has the largest MIN value, the portfolio allocation will put virtually everything in Gamble A as the risk tolerance tends towards 0. And since Gamble C has the largest EMV the allocation will put virtually everything in Gamble C as the risk tolerance tends towards infinity. With the RTP-MVF analyses, one gets the allocations for in between risk tolerance values which can be plotted to give a pictorial representation for the MVF solutions. This is presented in the two charts below.

The first chart shows the fractional asset allocations as risk tolerance varies from 0 to 100, which has sufficient resolution to see the details of the allocations as fa drops from 1 to 0 as shown in the Figures 10 and 11 (when the risk tolerance is about 21.5). Although the three allocations vary nonlinearly with risk tolerance on this interval, they also exhibit an approximately piecewise linear shape if we break the interval from 0 to 21.5 in two intervals, from 0 to 6 and from 6 to 21.5. From the second chart it is seen that once A is out of the allocation (risk tolerance > 21.5), the allocations to B and C vary in an almost perfectly linearly way with respect to risk tolerance, crossing with a 50-50 allocation when risk tolerance is about 40. We do not at present have an explanation for this behavior, but when it happens it is very convenient since the results of the analysis can be summarized in a breakpoint table in which one can interpolate to find approximately optimal allocations for all risk tolerance levels. For this example, one obtains the following table. This Table 11 is analogous to the Policy Region Table in the discrete case, but allows for fractional allocations to the investment options rather than the discrete selection in the earlier analysis.

Risk Tolerance	Fraction in A	Fraction in B	Fraction in C
0	1	0	0
6	0.32559	0.3932	0.28121
21.5	0	0.52338	0.47662
450	0	0	1

Table 11: Asset Allocation Table.

With the piecewise linear tabulation, one can very easily obtain solutions for situations like the following. Suppose an investor says "I don't really know what my risk tolerance is, but I feel that I should be making approximately equal investments in all three gambles." The analyst can then point out that the breakpoint table shows that approximately equal allocations occur for the risk tolerance of 6, and after a little searching with the nonlinear optimizer can announce that the closest approach to equal allocations on the maximal value frontier occurs for $\tau=6.2$ at which point the allocations are (.318766, .395279, .285955) respectively, or in dollars, (\$15.84, \$19.76, \$14.30). On the other hand, suppose an investor says "I feel that I should have at least as much in Gamble A than in the other two combined." The breakpoint table shows that in this case the risk tolerance would have to be less than 6 (in the first interval of the table) and again a little searching with the nonlinear optimizer shows that this relationship between the allocations occurs for $\tau=3.44$ at which point the allocations in dollars are (\$25.00, \$15.54, and \$9.46)

On the other hand, if the investor says “I would like to have twice as much invested in Gamble B than the other two combined,” the analyst has to respond that this does not occur on the maximal value frontier, and should be set aside as a goal since it is not optimal for any risk tolerance. Of course, a formal risk tolerance assessment will give more precise guidance regarding preferred allocations, but our point here is that valuable insights can be obtained even without a risk tolerance assessment, if one just carries out the requisite risk tolerance parametrics and tabulates the optimal solutions for a suitable range of risk tolerance values.

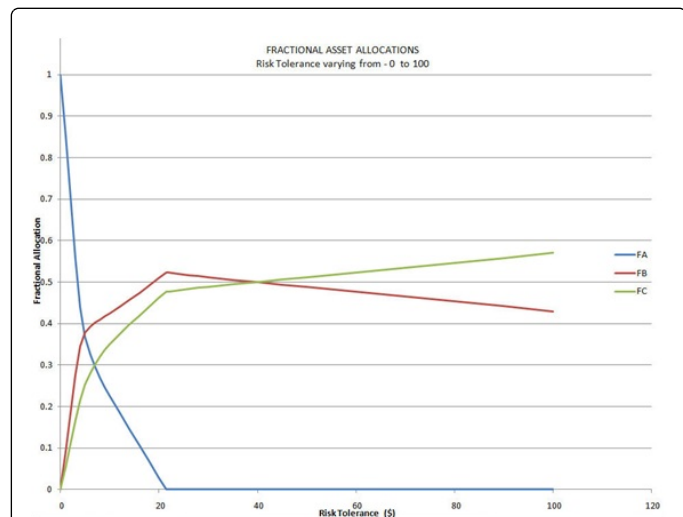


Figure 10: Asset Allocation as Risk Tolerance varies from 0 to 100.

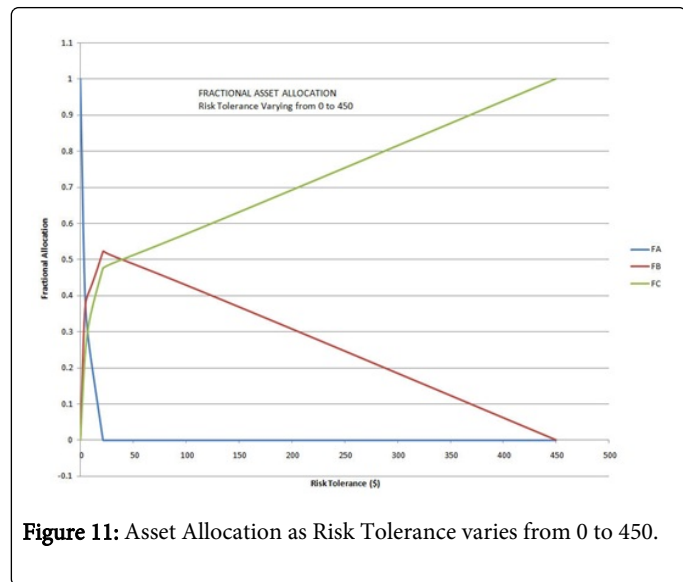


Figure 11: Asset Allocation as Risk Tolerance varies from 0 to 450.

Least Square Risk Tolerance Estimation

In the foregoing, we have been concerned with characterizing solution behavior as risk tolerance varies across its entire range from 0 to +infinity. Now we consider how to narrow the range of risk tolerances by making an assessment of the risk tolerance that best represents the choice behavior of the decision maker when faced with a systematically constructed sequence of “calibration gambles.” These

calibration gambles are kept simple by using 50-50 “flip of the coin” type lotteries with payoffs within the range of outcomes for the decision tree under consideration.

Generalized Interview Method to Estimate Risk Tolerance

While there is an objective risk associated with any particular gamble (risky option), the risk attitude that we measure by means of the risk tolerance is a subjective attitude embedded in the heart and mind of the decision maker (D.M.) faced with the choice. Hence to estimate risk tolerance, we must ask the D.M. to state some preferences about certain simplified “calibration” gambles. We then find a “best fit” to the revealed preferences of the D.M. and use the estimated tolerance to evaluate more complex gambles on behalf of the D.M. The method described here is a generalization and extension of the one presented in Chapter 8 of the Smart Choices text by Hammond et al. [11,12].

The estimation of risk tolerance is best done in relation to a particular decision problem that can be modeled by means of a decision tree. If we let H be the maximum payoff in the tree (for High) and L be the minimum payoff in the tree (for Low), then a sequence of structured calibration gambles can be developed in the following way. The notation for the certainty equivalents comes from imagining that we are construction a utility function for money where $U(L) = 0$ and $U(H) = 100$ (as in the Smart Choices text).

First Calibration Gamble (G.L-H): 50-50 chance for either L or H

The expected utility for this gamble is $(U(L)+U(H))/2$ or 50. We obtain the empirical (subjective) certainty equivalent for this gamble by asking the question: What is the least amount of cash in an envelope that you would find equivalent in preference to G.L-H? Or, What is the least amount of cash in an envelope that you would accept instead of receiving G.L-H? Or, what is the most amount of cash in an envelope that you would reject in preference for G.L-H? These are meant to be questions that all have the same answer, just asked in different ways. The answer will be called ECE.50 for Empirical Cash Equivalent for G.L-H.

Second Calibration Gamble (G.L-ECE.50): 50-50 chance for either L or ECE.50

The expected utility for this gamble is $(U(L)+U(ECE.50))/2$ or 25. The answer to the same type of question in this case will be called ECE.25 for Empirical Cash Equivalent for G.L-ECE.50.

Third Calibration Gamble (G.ECE.50-H): 50-50 chance for either ECE.50 or H

The expected utility for this gamble is $(U(ECE.50)+U(H))/2$ or 75. The answer to the same type of question in this case will be called ECE.75 for Empirical Cash Equivalent for G.ECE.50-H.

We could stop there, but in the Smart Choices case they do one more bisection in preference between ECE.75 and H.

Fourth Calibration Gamble (G.ECE.75-H): 50-50 chance for either ECE.75 and H

The expected utility for this gamble is $(U(ECE.75)+U(H))/2$ or 87.5. So the answer to the same type of question in this case will be called ECE.875 for Empirical Cash Equivalent for G.ECE.75-H.

Having these four subjective cash equivalent evaluations from the D.M. for the four calibration gambles, we now consider the question as to what risk tolerance best represents these preferences or preference equivalences that have been given by the D.M. We would like to have a range of imputed values as well as a “least square” estimate, this being the “best fit” in terms of minimizing the sum of squared deviations between the stated (empirical) CE’s and those that are imputed by the CE function for Exponential Utility. In fact, each of the empirical CEV responses of the decision maker can be used to obtain an estimate of risk tolerance, simply by solving for that risk tolerance which would give a theoretical CEV which matches the stated CEV exactly. Hence Solving

$$-\tau \ln(.5e^{-L/\tau} + .5e^{-H/\tau}) = ECE.50 \text{ yields } \tau.50$$

In the same way, one obtains three other risk tolerance estimates, as follows

$$\text{solving } -\tau \ln(.5e^{-L/\tau} + .5e^{-ECE.50/\tau}) = ECE.25 \text{ yields } \tau.25$$

$$\text{solving } -\tau \ln(.5e^{-ECE.50/\tau} + .5e^{-H/\tau}) = ECE.75 \text{ yields } \tau.75$$

$$\text{solving } -\tau \ln(.5e^{-ECE.75/\tau} + .5e^{-H/\tau}) = ECE.875 \text{ yields } \tau.875$$

These four estimates will most likely be all different (people are not naturally consistent, nor do they have a built in exponential utility function). Hence the least square estimate we find as described below will be somewhere in the interval between $\tau.min$ and $\tau.max$ where $\tau.min = \min(\tau.25, \tau.50, \tau.75, \tau.875)$ and $\tau.max = \max(\tau.25, \tau.50, \tau.75, \tau.875)$. In this paper we refer to the interval $[\tau.min, \tau.max]$ as the “interval of uncertainty” for the risk tolerance estimate for the decision maker in question.

To obtain our “optimal” estimate we invoke the time honored principle of least square fits. In this case, we want to minimize the root mean square of the deviations between the observed and the fitted certainty equivalents. Hence our maximand is

$$RMS = [(ECE.50 + \tau \ln(.5e^{-H/\tau} + .5e^{-L/\tau}))^2 + (ECE.25 + \tau \ln(.5e^{-ECE.50/\tau} + .5e^{-L/\tau}))^2 + (ECE.75 + \tau \ln(.5e^{-H/\tau} + .5e^{-ECE.50/\tau}))^2 + (ECE.875 + \tau \ln(.5e^{-H/\tau} + .5e^{-ECE.75/\tau}))^2 / 4]^{1/2}$$

This RMS deviation measure is minimized with respect to variations in the risk tolerance τ , easily accomplished with the Excel Solver, for example. The risk tolerance that minimizes the sum of squared deviations is said to be the least square estimate of the decision maker’s risk tolerance. This will lie in the interval of uncertainty, which is between $\tau.min$ and $\tau.max$, but not necessarily half way in between the two end points.

Conclusions

We have presented a general risk tolerance parametric analysis procedure (RTP) that can be applied to any decision making problem under uncertainty with risk-averse preferences. This methodology produces two principal results: a maximum value function of risk tolerance (the MVF) and an optimal policy region tabulation (PRT), or asset allocation breakpoint table that describes all of the solutions that occur as risk tolerance is varied from (close to) 0 to very large values.

This is seen as a preliminary model validation method that is done early on in the search for a “requisite decision model” for the situation at hand. It can be done prior to risk tolerance estimation for the decision maker in question, so that when the risk tolerance estimate is developed, the optimal solution(s) for the interval of uncertainty around the estimate can be read directly from the Policy Region Tabulation or the Asset Allocation Breakpoint Table that have previously been determined.

In addition, for decisions involving an information acquisition option, the plot of CEVSI may indicate information acquisition over a significant range of risk tolerance values where the EVSI results do not. This information acquisition can mitigate downside risk significantly which manifests itself in greater CEVSI that is not reflected in the EVSI measure of value.

It is believed that this risk tolerance parametrics methodology will lead to better models, greater insight and less risk exposure in the analysis and optimization of decisions made under uncertainty with risk-averse preferences.

References

1. Raiffa H (1968) *Decision Analysis*, Reading MA: Addison-Wesley.
2. Howard RA (1980) “Risk Preference,” paper 34 in *The Principles and Applications of Decision Analysis*, Vol II, Edited by Ronald A. Howard and James E. Matheson, Strategic Decisions Group, Menlo Park, California, April 1984.
3. Davis RE (2001) “Decision Policy Optimization Via Certainty Equivalent Functions for Exponential Utility” in Chapter 6, Vol. 6 of *Advances in Mathematical Programming and Financial Planning* (Elsevier Science Ltd.), pp. 89-113.
4. Davis RE (2002) “Portfolio Optimization with beta distributed returns and exponential utility” presented at the Hawaii Int’l Conference on Statistics & Related Fields, Honolulu, HI.
5. Davis RE (1993) “Certainty Equivalent Models for Portfolio Optimization using Exponential Utility and Gamma Distributed Returns” in Chapter 4, Vol. 3 of *Advances in Mathematical Programming and Financial Planning* (Greenwich, CT: JAI Press), (Published December 1992) pp. 69-108.
6. Freund RJ (1956) *The Introduction of Risk into a Programming Model*. *Econometrica* 38: 410-421.
7. Billingsley P (1995) *Probability and Measure*, NY NY: John Wiley & Sons, Italy.
8. Clemen R, Reilly T (2001) *Making Hard Decisions*, with DecisionTools, Pacific Grove, CA, Duxbury Press, Italy.
9. Davis RE (1995) “An Empirical Back Test of the Portfolio Gamma Model for Optimal Asset Allocation” in Chapter 5, Vol. 4 of *Advances in Mathematical Programming and Financial Planning* (Greenwich, CT: JAI Press), (Published Dec. 1995), pp. 77-101.
10. Fox C (2012) “Creating Value with RISK Appetite & RISK Tolerance Statements” in *Financial Executive*. 9: 93-95.
11. Hammond J, Keeney R, Raiffa H (1999) *SMART CHOICES: A Practical Guide to Making Better Life Decisions*, New York, Broadway Books, USA.
12. Pratt J, Raiffa H, Schlaifer R (1995) *Introduction to Statistical Decision Theory*, Cambridge, MA, MIT Press, UK.