CEE 631
CONSTRUCTION DECISIONS UNDER UNCERTAINTY

Class Notes

INTRODUCTION TO DECISION ANALYSIS



© 2024 Photios G. Ioannou, Ph.D., Dipl.Eng., F.ASCE
Professor of Civil and Environmental Engineering
Construction Engineering & Management
University of Michigan
Ann Arbor, Michigan 48109-2125

Decision Trees

Construction projects are dynamic, changing from one day to another. The one constant seems to be change. Management of many projects, if not most, consists primarily of meeting immediate needs with little attention for the future. This atmosphere of "putting out fires" can be characterized as a sequence of decisions under uncertainty and their resulting outcomes.

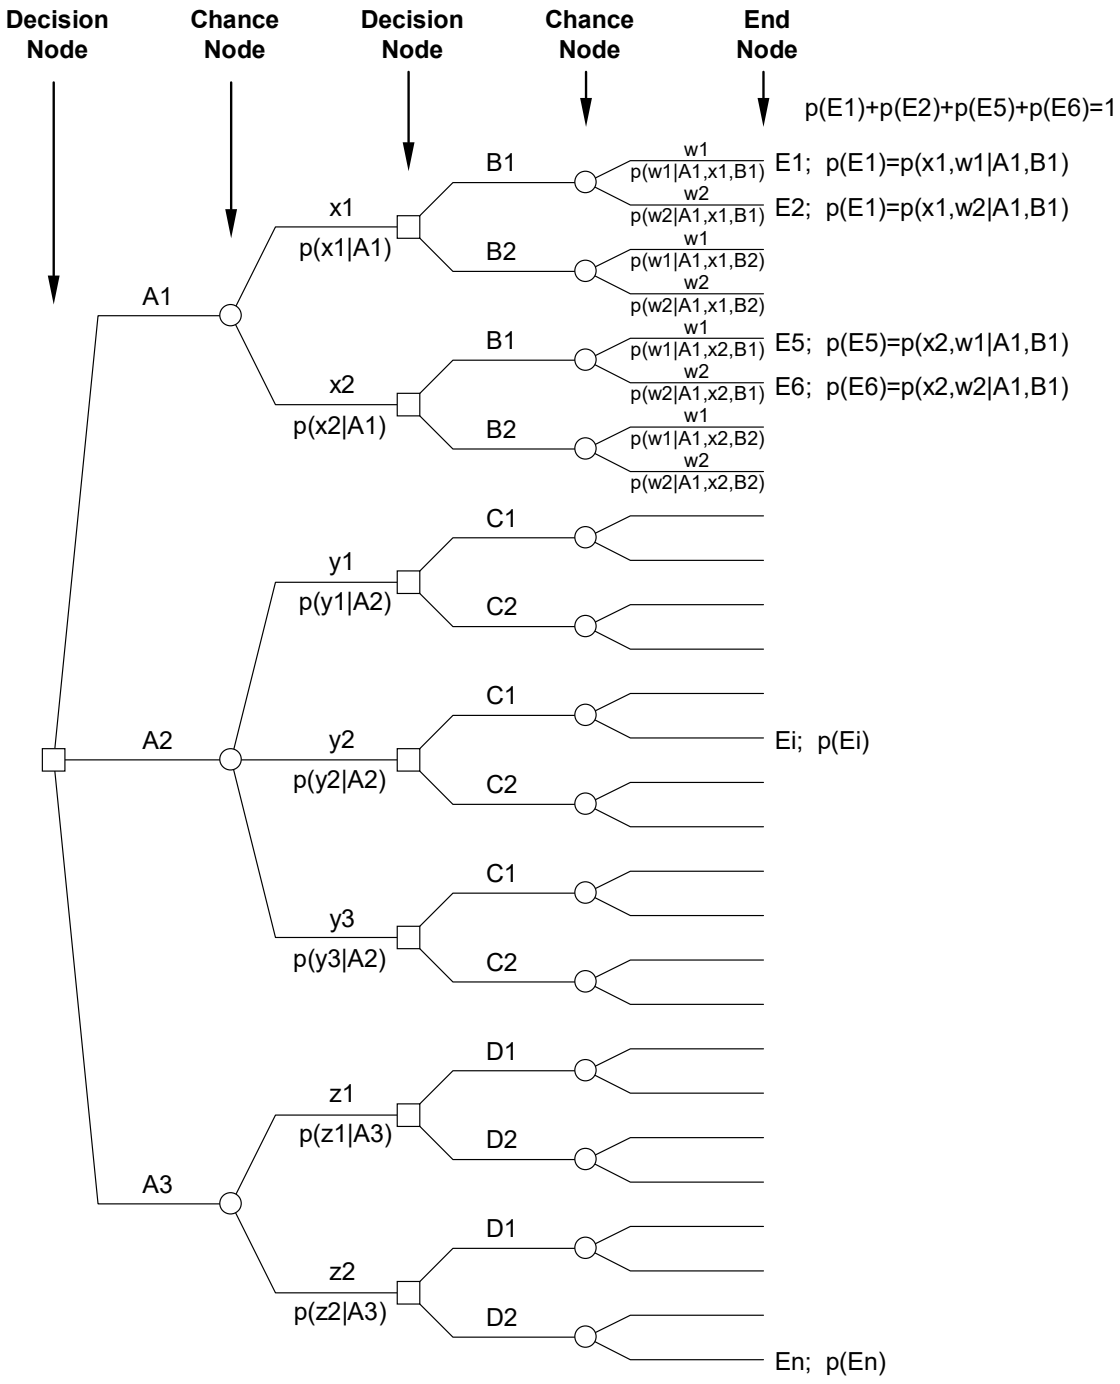
Each manager intuitively understands that in this sequence, each decision not only affects its immediate outcome, but also affects future decisions. Patterns of management are built around the desire to keep one's options open. When a decision is made on bidding a project, a contractor realizes that he must consider the important decisions to be made if he gets the contract. A designer leaves flexibility in a facility for changes (decisions) after the facility is in use and its actual utilization has been determined. When a soils engineer makes a decision concerning the design of an exploration program, he carries in his mind the design decisions that will be made later, after the samples are known.

A decision sequence can be diagrammed using a *decision tree*. Using a tree representation makes the decision problem easier to understand, analyze, and communicate to others. A decision tree consists of nodes and branches as shown in Fig. E.1. There are three types of nodes in a decision tree: decision nodes (squares), chance nodes (circles) and end-nodes or value-nodes (triangles). A decision node is shown using a small square whose outgoing branches to the right represent the possible decision alternatives (choices). When faced with that particular decision, a decision maker must select one of the alternatives represented by the branches immediately to the right of the decision node. A chance node is similar to a random variable and is shown as a small circle. All the possible outcomes of a chance node (i.e., the sample space of the random variable) are represented by the outgoing branches to its right. Thus, when a chance node materializes in real life, exactly one of the outgoing branches must occur. This means that the branches emanating from a chance node represent mutually exclusive and collectively exhaustive events that define the complete sample space for the underlying random variable. Hence, the probabilities associated with the branches of a chance node must always add up to 100%.

In Fig. E.1, the first decision has alternatives A_1, A_2, and A_3 of which one must be selected. If A_1 is selected, then one of two possible chance events, x_1 and x_2, will occur, each with its own probability of occurring. Since x_1 and x_2 represent all possible outcomes of that chance node, $P[x_1]+P[x_2] = 1$. If A_2 were selected instead of A_1, then there are three possible outcomes y_1, y_2, y_3 with the probabilities shown. If A_1 were chosen and outcome x_1 resulted, then the decision maker would have to make a decision between alternatives B_1 and B_2, each of which would have its own possible uncertain outcomes.

At each decision node one alternative must be chosen, and at each chance node one event will occur. A decision node may be followed by another, but these could be resolved into a single decision. For example, a decision to bid could be followed by a markup decision with alternatives of 5% and 7%. However, the two decisions could be consolidated to one with alternatives (1) bid at 5%, (2) bid at 7%, or (3) do not bid. Likewise, a chance node could be followed by another; however, they could be put together as the intersection of their events.

© 2024 *Photios G. Ioannou*



Structure of a Decision Tree

Branches at a Chance Node = Feasible events *given* the preceding Decision and Chance branches.
Chance Node branches represent mutually exclusive (ME) and collectively exhaustive (CE) events.
Probabilities of Chance Node branches are *conditiona*l on all preceding Decision and Chance branches.
Sum of probabilities assigned to the branches of a Chance Node = 1.0
Expected Monetary Value Decision Rule (EMV): At any Decision Node select the alternative that leads to the *greatest expected profit* (or *smallest expected cost*).

Figure E.1. Structure of a Decision Tree

Terminology: Decisions – Alternatives – Distinctions –Possibilities

A decision involves making a *choice* from a set of *alternatives* where each alternative leads to a possible *future*. Selecting the best decision alternative requires the evaluation of the futures associated with each choice. The description of the *future* for each decision alternative is called *characterization*. Each decision alternative is *characterized* by defining the *distinctions* (chance variables, etc.) and *possibilities* (chance outcomes, etc.) that describe the associated future.

In the context of decision analysis, *uncontrollable* variables, are those variables whose outcomes are not controlled directly by the decision maker. The outcomes of *uncontrollable* variables are determined by chance (or, equivalently by other decision makers).

Uncontrollable variables can be described by several technical terms, such as

- *Distinctions*
- *Chance variables*
- *State variables*
- *Random variables, etc.*

A *distinction* describes the possible future outcomes of a process that is not directly controlled by the decision maker. For example, the *wind speed* at a specific construction site at noon tomorrow is a *distinction*.

The outcomes of a distinction are called *possibilities*. A flip of a coin, for example, is a *simple distinction* that has only two *possibilities* (*heads* or *tails*).

The main reason for this relatively new terminology (i.e., *distinctions*, *possibilities*) is to stress that in decision analysis, distinctions and possibilities are often defined in a way that suits the decision problem at hand.

For example, in general, wind speed can be modeled as follows:

- *Wind speed* is a continuous random variable with an infinite sample space over all positive real numbers.

In a decision problem, however, where a construction operation can take place only if wind speed is less than 5 mph, we may wish to model wind speed as follows:

- *Wind speed* is a simple distinction (i.e., as a discrete binary *Bernoulli* random variable) with two possibilities (possible outcomes of interest)
  - *Favorable* (≤ 5mph), and
  - *Unfavorable* (> 5mph).

© 2024 *Photios G. Ioannou*

Definition Requirements for Distinctions and Possibilities

In decision analysis, it is often tempting to define distinctions and the associated possibilities using *loose* (imprecise) terms, such as "*demand* will be one of *high*, *medium*, *low*". Notice that these terms, unless defined clearly, can mean different things to different people. Clearly, such ambiguity must be avoided. Thus, for distinctions and possibilities to be defined correctly, they must pass the following three tests:

- *Clarity*
- *Observability*
- *Usefulness*

These three tests are further described below.

The Clarity Test

Distinctions and possibilities need to be defined clearly (accurately) so that they pass the *clarity test*.

Example: The following distinction and possibilities pass the clarity test:

- The *Demand* (*D*) *distinction* represents:
  - "Units o*rdered for concrete precast tee sections, model no. 558, in the 1st quarter of 2015*".
- The three *demand possibilities* are defined as:
  - *Low*: $D \leq 250$ units
  - *Medium*: 250 units $< D \leq 500$ units
  - *High*: $D > 500$ units

An effective way to explain the *clarity test* to others (e.g., to experts within an organization) is to present it as a *clairvoyant test*. In this context, a *clairvoyant* is a (fictitious) person who can state infallibly or predict accurately the outcome of any past or future distinction.

The clairvoyant test:

> For a distinction to pass the clarity test it must be specified sufficiently to allow a clairvoyant to state, *without any exercise of judgment*, which particular possibility for that distinction has occurred (if in the past) or will occur (if in the future).

For example, a contractor may be interested to know if the weather next Monday will be "favorable". One could argue that this is a *simple distinction*. The weather will be "favorable" or "not favorable". However, the distinction "the weather will be favorable next Monday" does not pass the clarity test. For a clairvoyant to be able to make a prediction, we must define the exact conditions of temperature, precipitation, hours of sunshine, wind speed, etc., *that are important to the contractor* and which make up *his* definition of "favorable" weather. Clearly, the definition of "favorable weather" would be very different for a painting contractor working outdoors and a farmer hoping for rain.

© 2024 *Photios G. Ioannou*

The Observability Test

A distinction is *observable* if the decision maker (a) already knows its outcome or (b) will know its outcome in the future. It is common sense that distinctions must be *observable* if future decisions will *depend* on them. This is clearly the case in gambling where only observable distinctions can create well-defined bets. For example, the toss of a coin is an observable distinction with two possibilities, "Heads" or "Tails". The decision who wins a bet and how much money to change hands will be made after the outcome of a toss is observed.

Examples of engineering distinctions that are particularly important to be *observable* are experiments, or tests, and their outcomes. Clearly, the *outcomes of experiments or tests* must be observable so that future decisions can be made *knowing* the observed outcome of the test.

However, *not* all distinctions in a decision problem *must* be observable. For example, "the price that a competitor is paying for structural steel" may never be known to us and yet that distinction may play an important role in our decision making. The fact that the outcomes of some distinctions cannot be observed does not exclude them from decision analysis.

The *technical specifications for construction projects* are an excellent example where engineers strive for both clarity and observability of distinctions. That is why technical specs often reference highly technical standards or definitions that have withstood the test of time. For example, it does little good to specify that concrete should be "of high quality". "High quality" needs to be adequately specified in terms of *clear* and *observable* distinctions (such as the *required slump* or the *compressive strength of cylinders*). I.e., *observability* is a key requirement for distinctions that describe *the outcome of tests* that will be used as a *basis* for *decisions*, such as whether to accept or reject the concrete.

The Usefulness Test (The "Fit for the Purpose" Test)

Distinctions should be defined so they are useful towards the objective of achieving clarity of action. There are many clear and even observable distinctions that could be defined that may be irrelevant for the specific decision at hand. Thus, they provide no benefit towards achieving decision clarity for the particular decision problem. For example, a distinction about the specific weather conditions this Monday would not very useful for an investment decision involving land development. However, it may very useful in deciding whether to pour concrete on that day.

The requirement of *usefulness* is not just a question of whether a distinction is *relevant* to the decision at hand. It also has to do with whether the distinction and its outcomes are the most *appropriate* for the situation. For example, geological investigation could include geophysical in-situ tests (seismic, resistive, etc.), specimen extraction (coring), lab tests, etc. In this context, *usefulness* requires that tests should be selected so as to best support the type of design or construction decision that needs to be made after the test outcomes are known.

Example – The Plant Expansion Decision Problem

Mr. Jackson is investigating setting up a plant to prefabricate wood trusses and other prefabricated wood modules for his use and sale to others. He needs to decide (a) whether or not to invest and (b) whether the better immediate investment is in a "large" or a "small" plant.

He expects that business will grow for about 2 or 3 years and that the amount of demand over the first two years is a good indicator of the demand for the future. The decision tree for this investment is shown in Fig. E.2. The first square is the immediate decision which is to select the size of the plant. What follows is the demand over the next 2 years, which is a matter of chance. If a small plant is built, the decision to expand or not will be made based on the experience of the first 2 years. Then the demand over the next 8 years is a matter of chance. A 10-year horizon (2 years + 8 years) is considered enough, and the analysis stops there.

Estimates are made of the cost of constructing a large plant now, a small plant now, and expansion of the small plant in 2 years from now. Estimates of the income from the plant for high and low demands are also made. All cashflow estimates are shown in Fig. E.3.

For example, a small plant will cost \$300,000 now, have maintenance costs of \$25,000/year, and can be expanded for \$350,000 with an additional \$25,000/year maintenance. A high demand over the next two years will yield income of \$60,000 the first year and \$120,000 the second. High demand over the next 8 years will yield \$240,000/year, unless the plant is small in which case income is limited by capacity to \$120,000/year.

Because the investment and return are over a 10-year period, the time-value of money is important. Fig. E.4 shows the present value of each cash flow shown on Fig. E.3, at a rate of return r = 10%.

For example, the present worth of the construction of the large plant is:

$$PV_L = -500^k - 50^k\,(P|A,10\%,10) + 150^k\,(P|F,\,10\%,10) = -\$749.4^k$$

and a high demand over the second 8 years will yield

$$PV_{H2} = 240^k\,(P|A,10\%,8)\,(P|F,10\%,2) = \$1,058.2^k$$

The decision tree for the prefabrication plant has a total of 12 possible end-node outcomes that are the possible combinations of decision alternatives and chance outcomes. Each of these has a present value to be computed.

For example, if a small plant is built (PV = -418.9^k), the demand in the first two years is high (PV = $+153.7^k$), a decision is made to expand the plant (PV = -359^k) and a high demand results in the next 8 years (PV = $+1058.2^k$), then the PV of that outcome is the sum of the PV of its elements:

$$PV = -418.9 + 153.7 - 359 + 1058.2 = +434.0^k$$

The PW of each outcome combination is shown as an end-node present value in Fig. E.5.

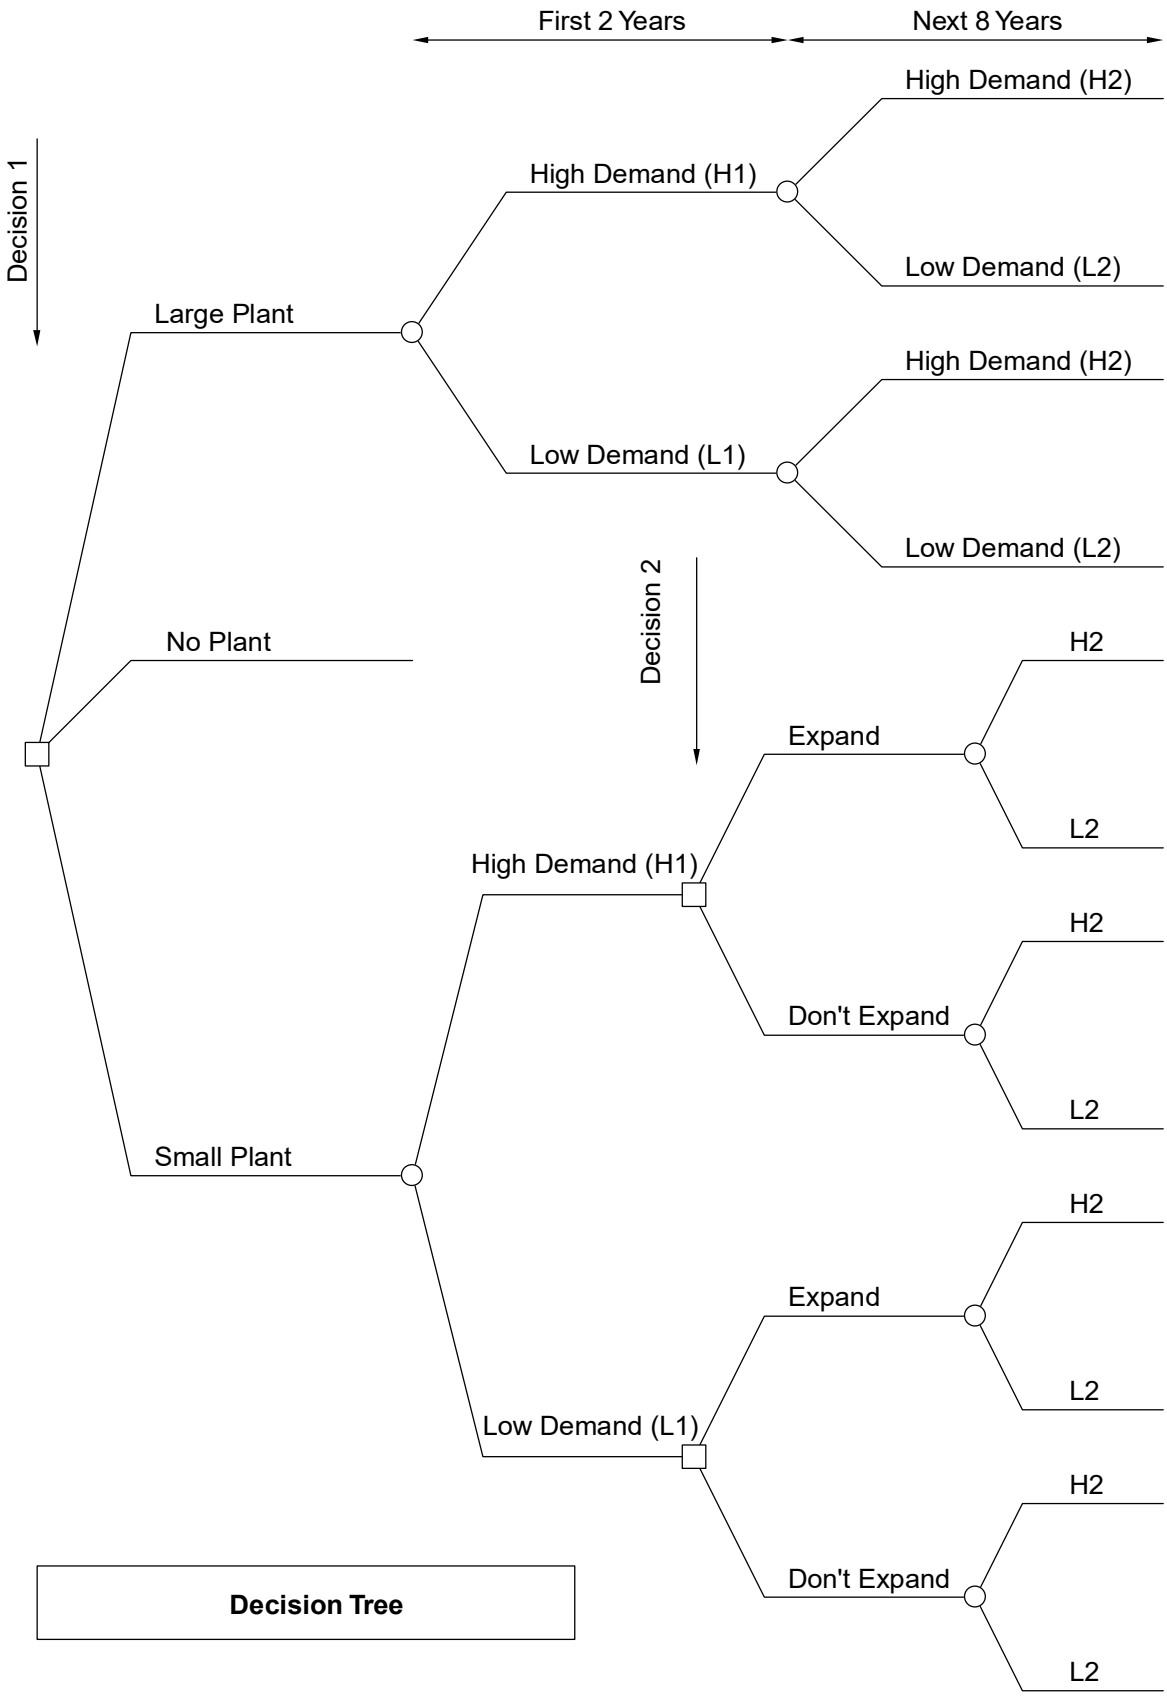


Figure E.2. Decision Tree for Plant Example

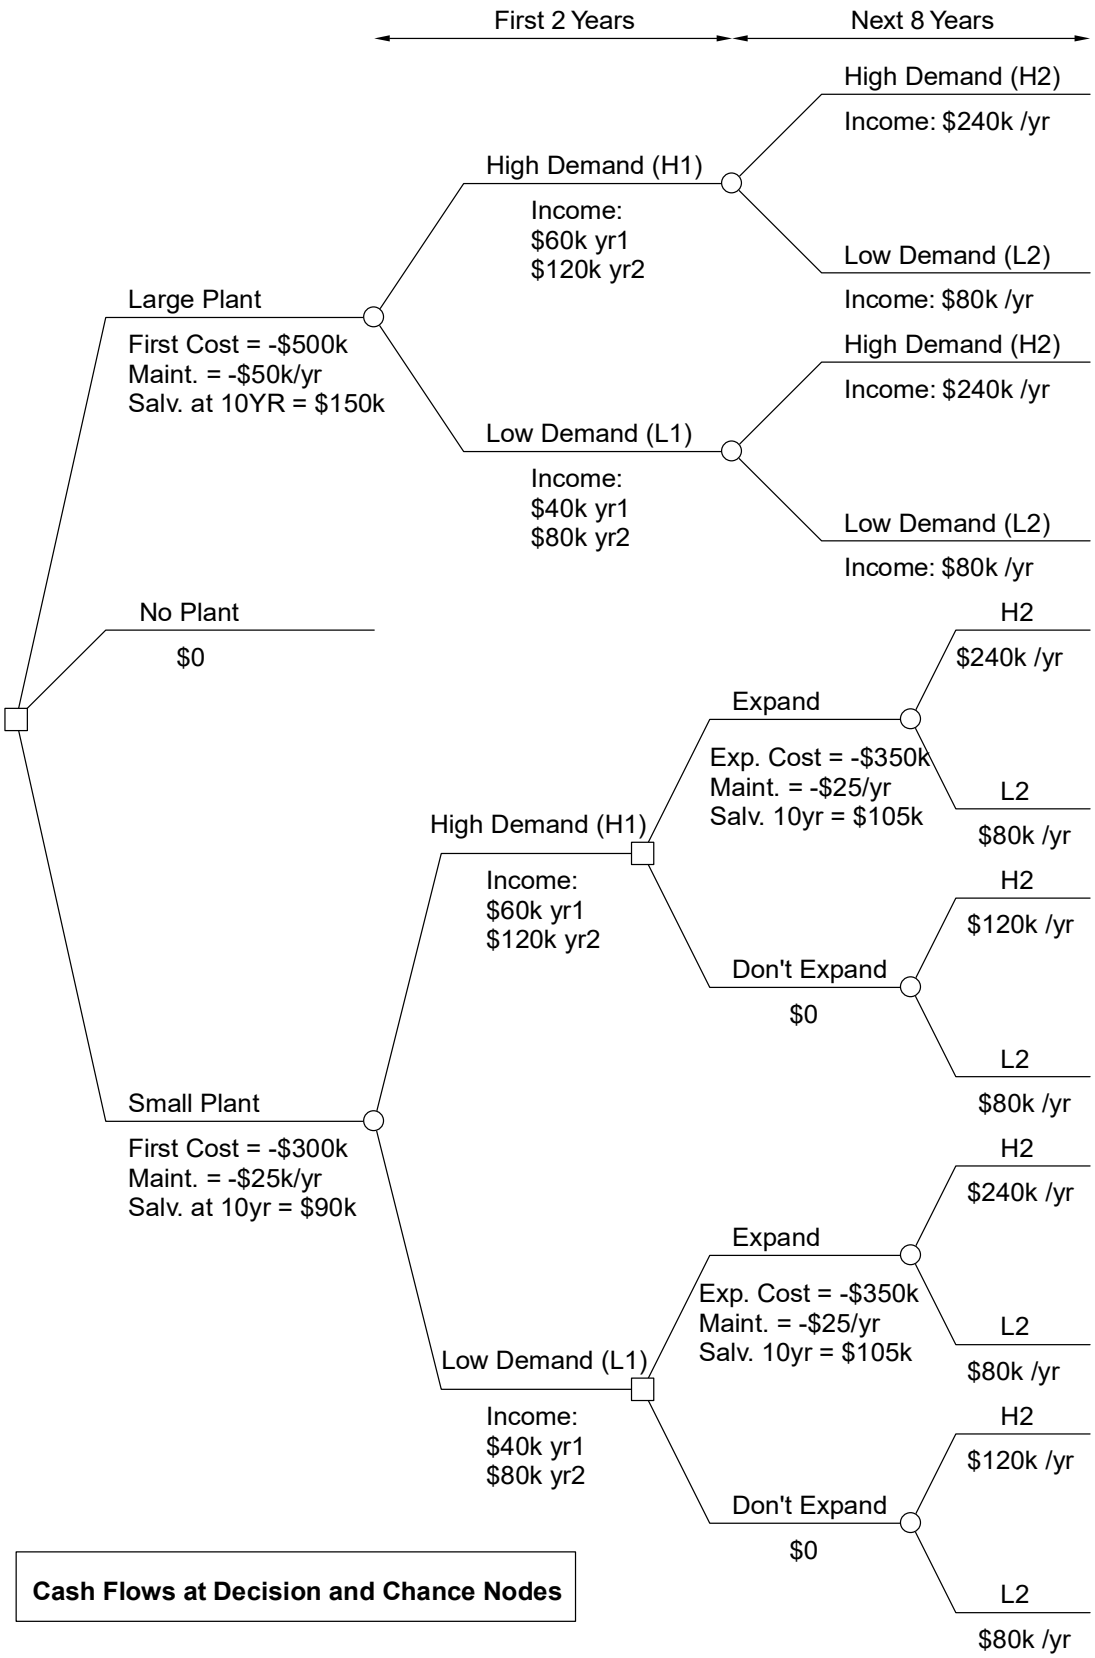


Figure E.3. Branch Cash Flows

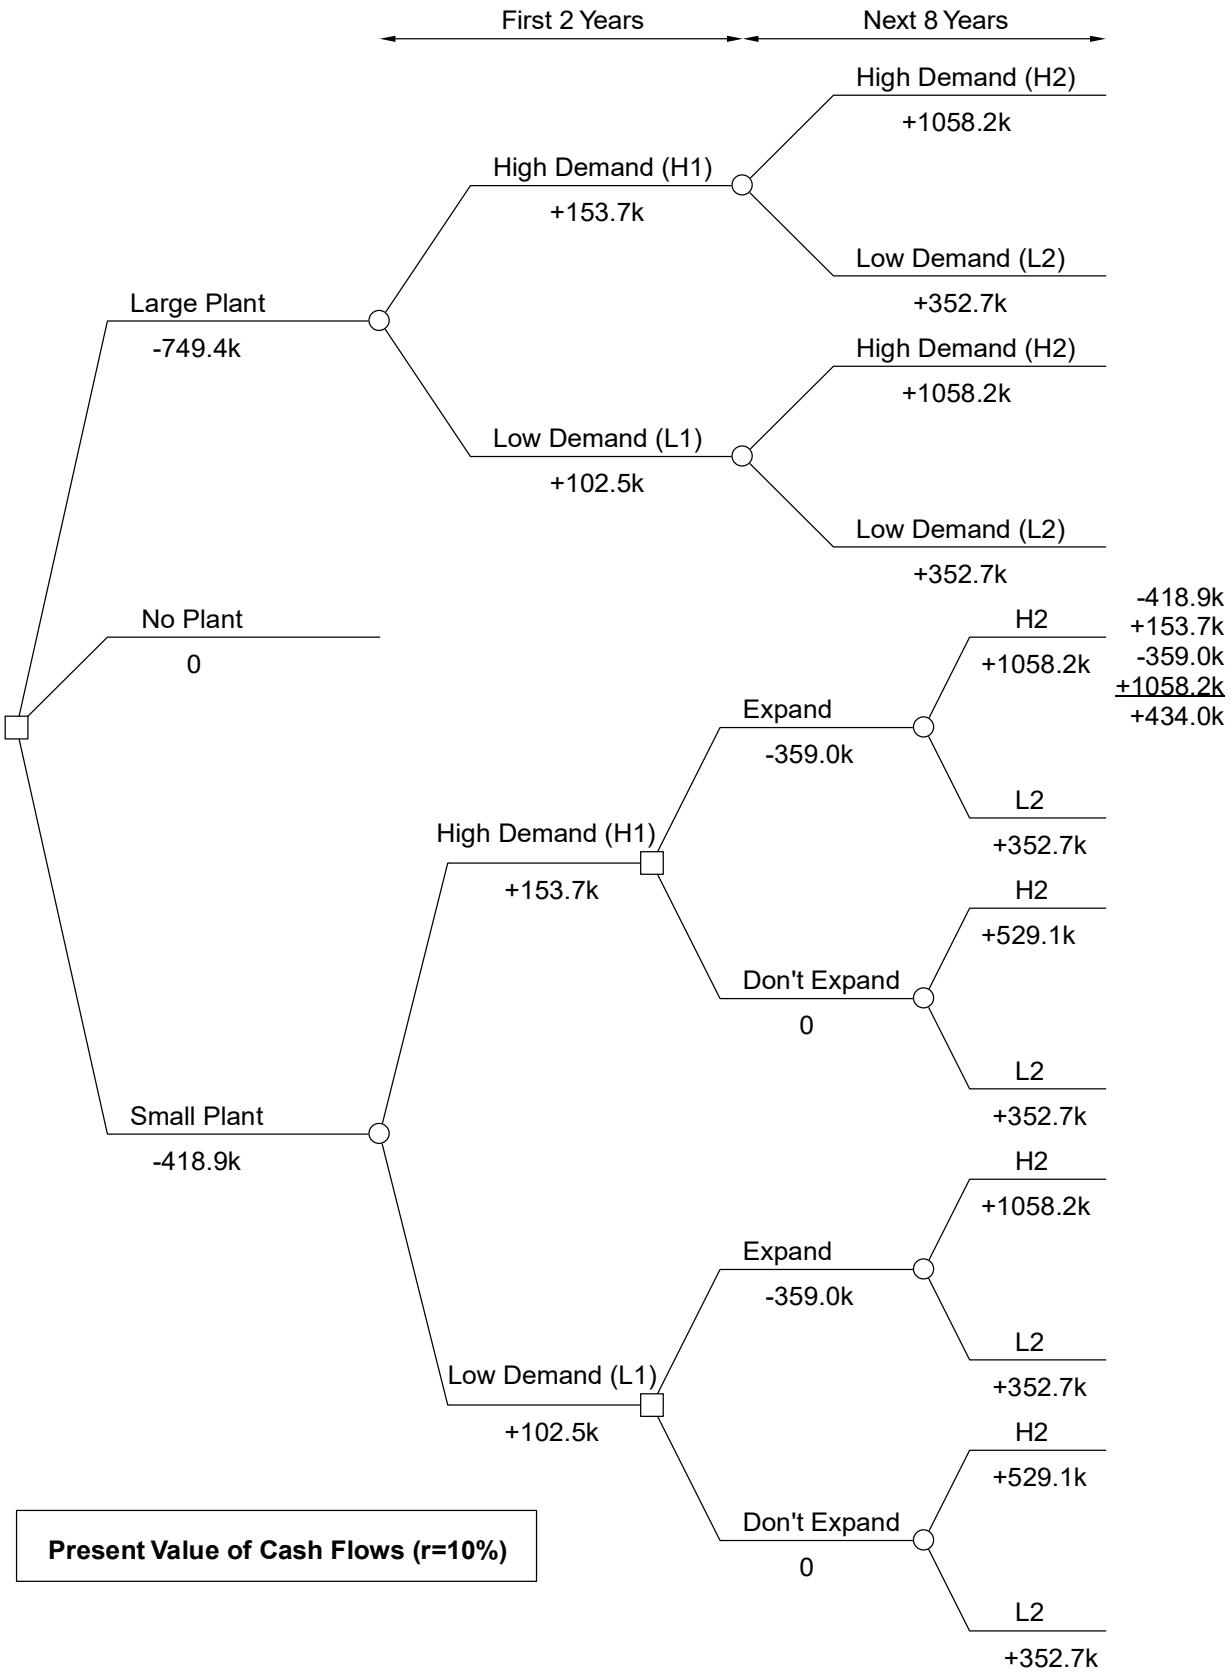


Figure E.4. Present Value of Branch Cash Flows

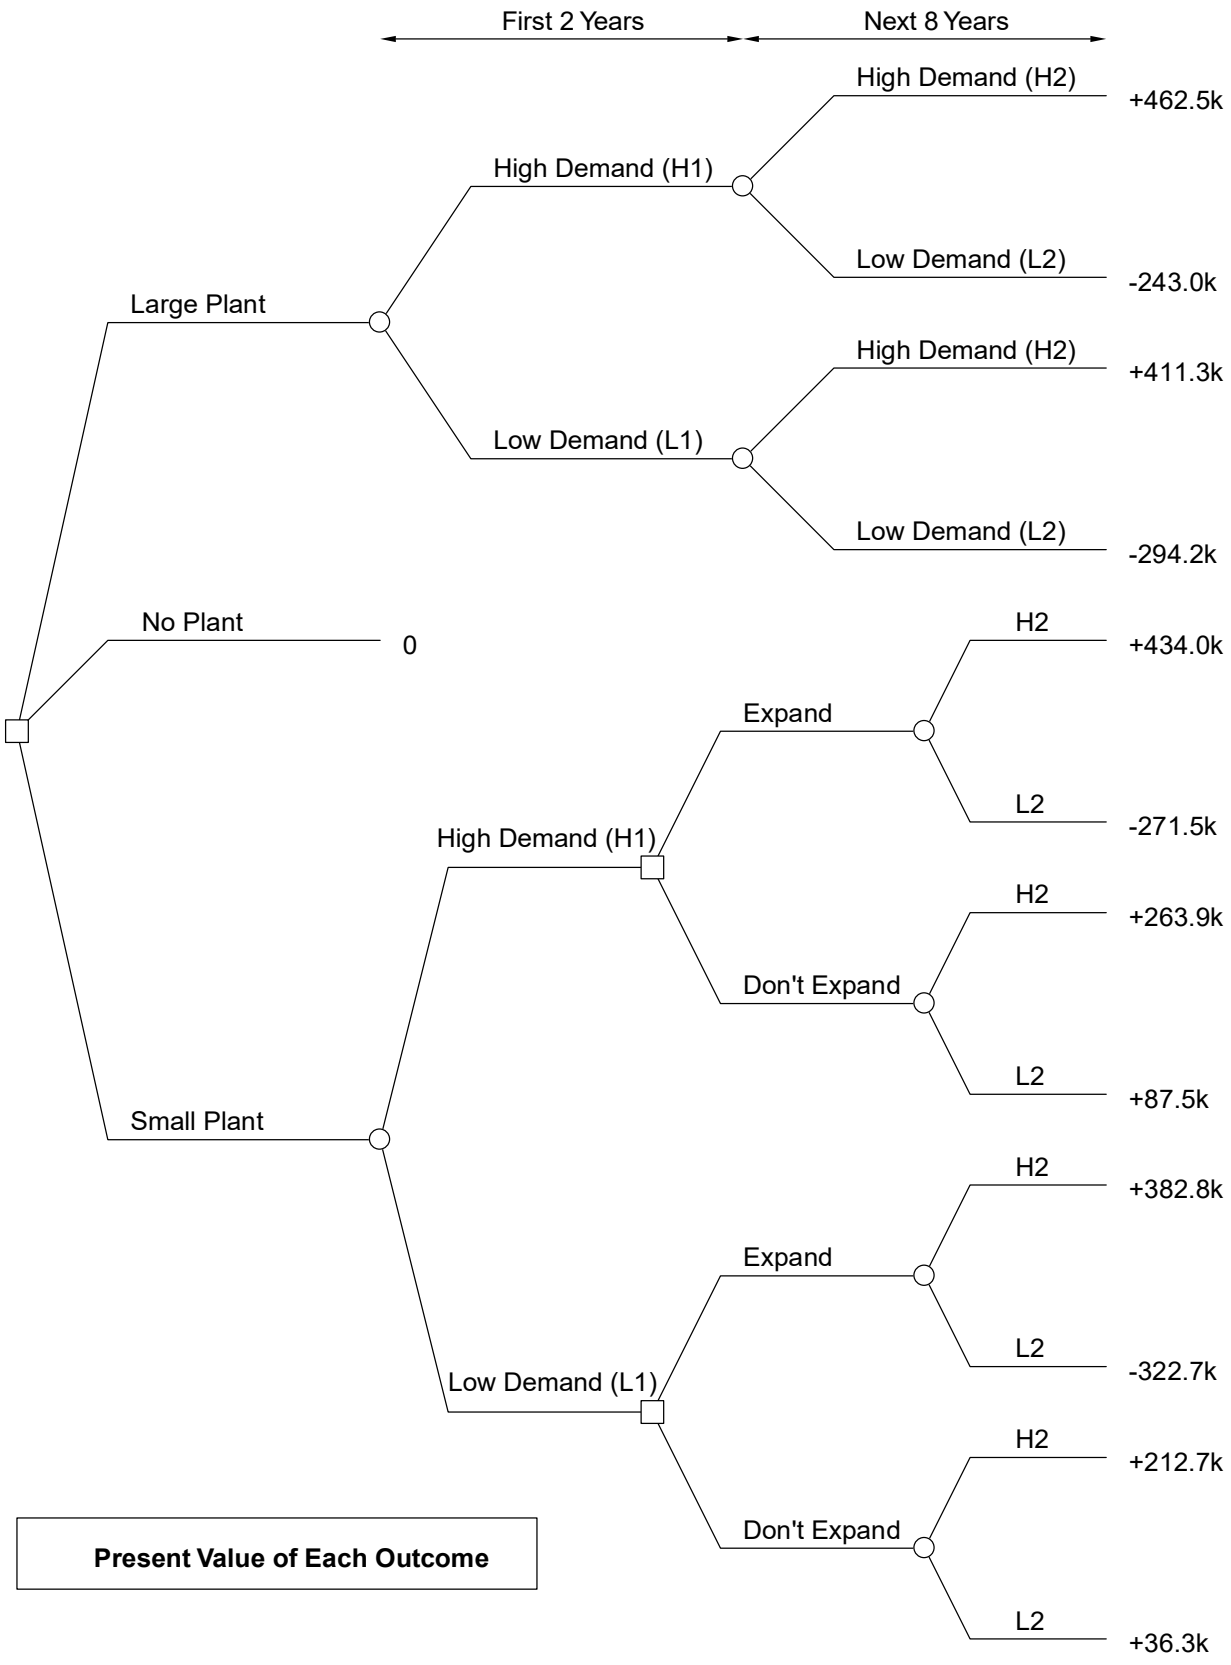


Figure E.5. End-Node Present Values

Nature's Tree

Decisions among alternatives are selected according to their expected values. The calculation of expected value requires an assignment of probabilities to each possible outcome on the decision tree. These must be estimated before a decision can be evaluated. It is helpful to show these probabilities on a "chance tree" or "nature's tree". For the prefabrication example, the probabilities estimated are shown on the nature's tree of Fig. E.6. We see that the probabilities for demand in the second period are conditional upon demand in the first.

We can construct a nature's tree for any events or conditional events, if we have the probabilities of their intersection (i.e., their joint probabilities). Joint probabilities can also be calculated from conditional probabilities as shown in Fig. E.6 on the right-hand side.

Rollback

Decision alternatives are evaluated by their expected values. We start with the present value of the end-node outcomes at the right side of the tree and move to the left by

(1) Computing the expected value of each chance node and
(2) Selecting the alternative with the highest expected value at each decision node.

This process is called "rolling back the tree" and is shown in Fig. E.7.

Based on the rollback in Fig. E.7, we see that the best decision is to

- Construct a small plant now because its expected value, E[V]=l55.l^{k}, is higher than that of the large plant, E[V]=l33.9^{k}.
    - If the initial demand over the first two years is high, then the plant should be expanded at that time, i.e., in two years.
    - If the demand over the first two years is low, then the plant should not be expanded.

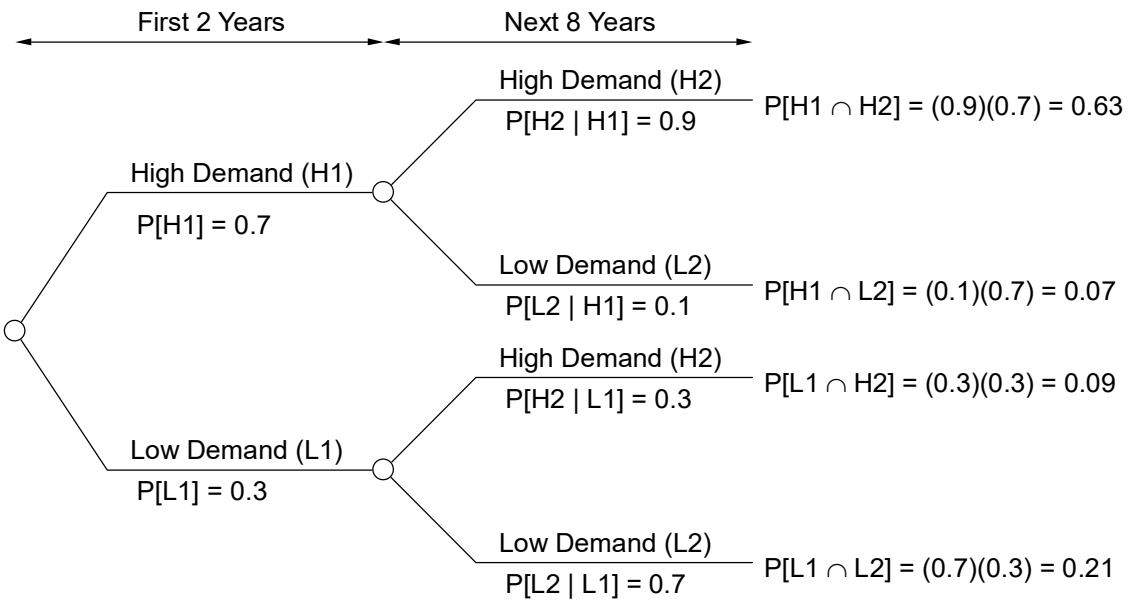
This constitutes the *optimal policy*, which tells us what to do at each decision node in the tree now and in the future. The decision branches for the optimal policy are shown in bold. The inferior decision branches are shown cut off with double cross lines.

The "Demand" Distinctions — Clarity and Observability

In this introductory example, the possibilities for the distinctions "Demand" in the first 2 years and the subsequent 8 years have been simplified to just two values "High" and "Low" that have been assumed to be defined in terms of revenue or units sold so that they are observable and pass the clarity test. Demand possibilities could also have multiple outcomes or be continuous.

A Simplified Example for the Decision to Purchase or Rent Equipment

Another simplified example that illustrates the decision to purchase or rent equipment for future work is shown starting with Fig. E.8. Fig. E.8 to E.10 illustrate how the analysis could be done by hand using pencil and paper. Fig. E.11 and E.12 show how the same example could be modeled using the Windows application "DPL" (Decision Programming Language).



Nature's Tree (Tree of Chance)

Assumption: The Demand probabilities are independent of the decision to build a large or small plant. This assumption is reasonable for this problem. In general, however, the probabilities at the various chance nodes depend on all preceding decision and chance branches. For example, the probabilities of technical success in an R&D effort depend on the decisions concerning the research methodology to be adopted.

Figure E.6. Probabilities on Nature's Tree

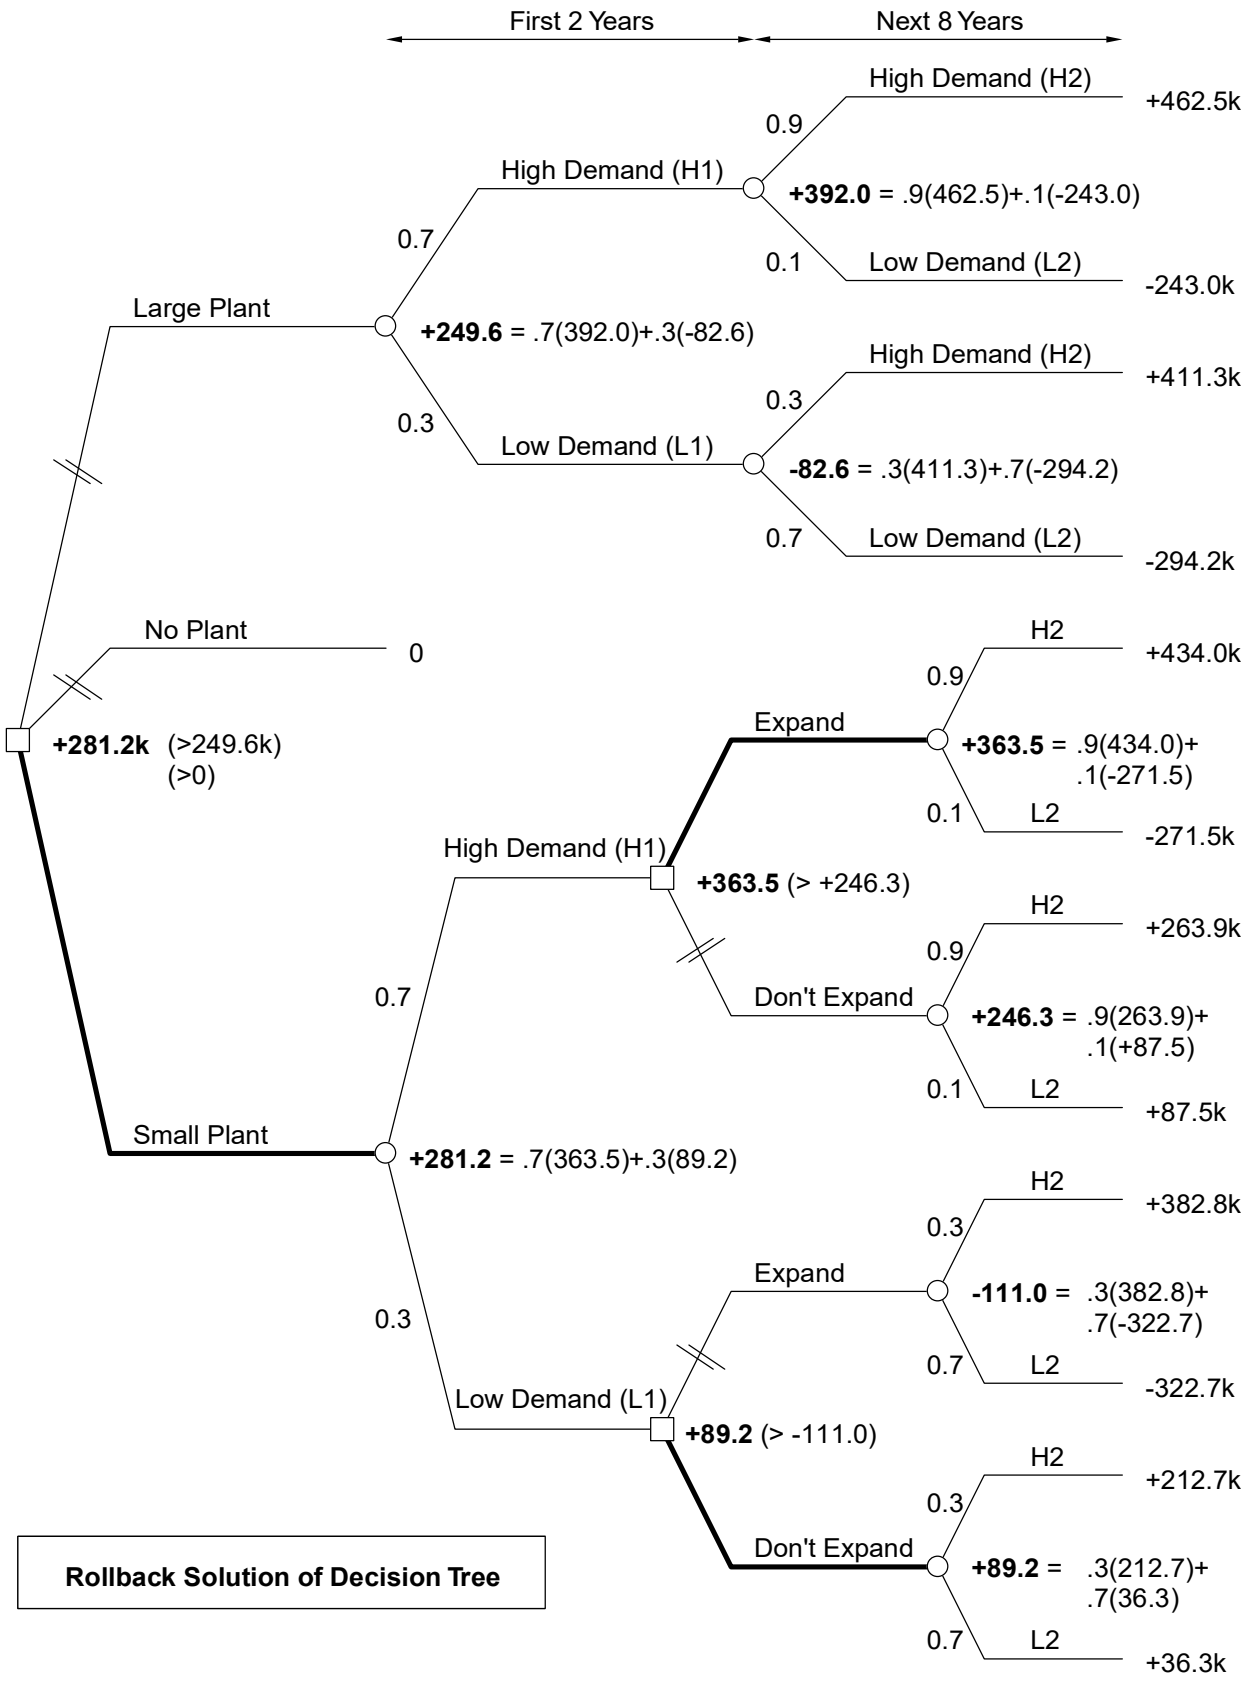


Figure E.7. Rollback of Decision Tree

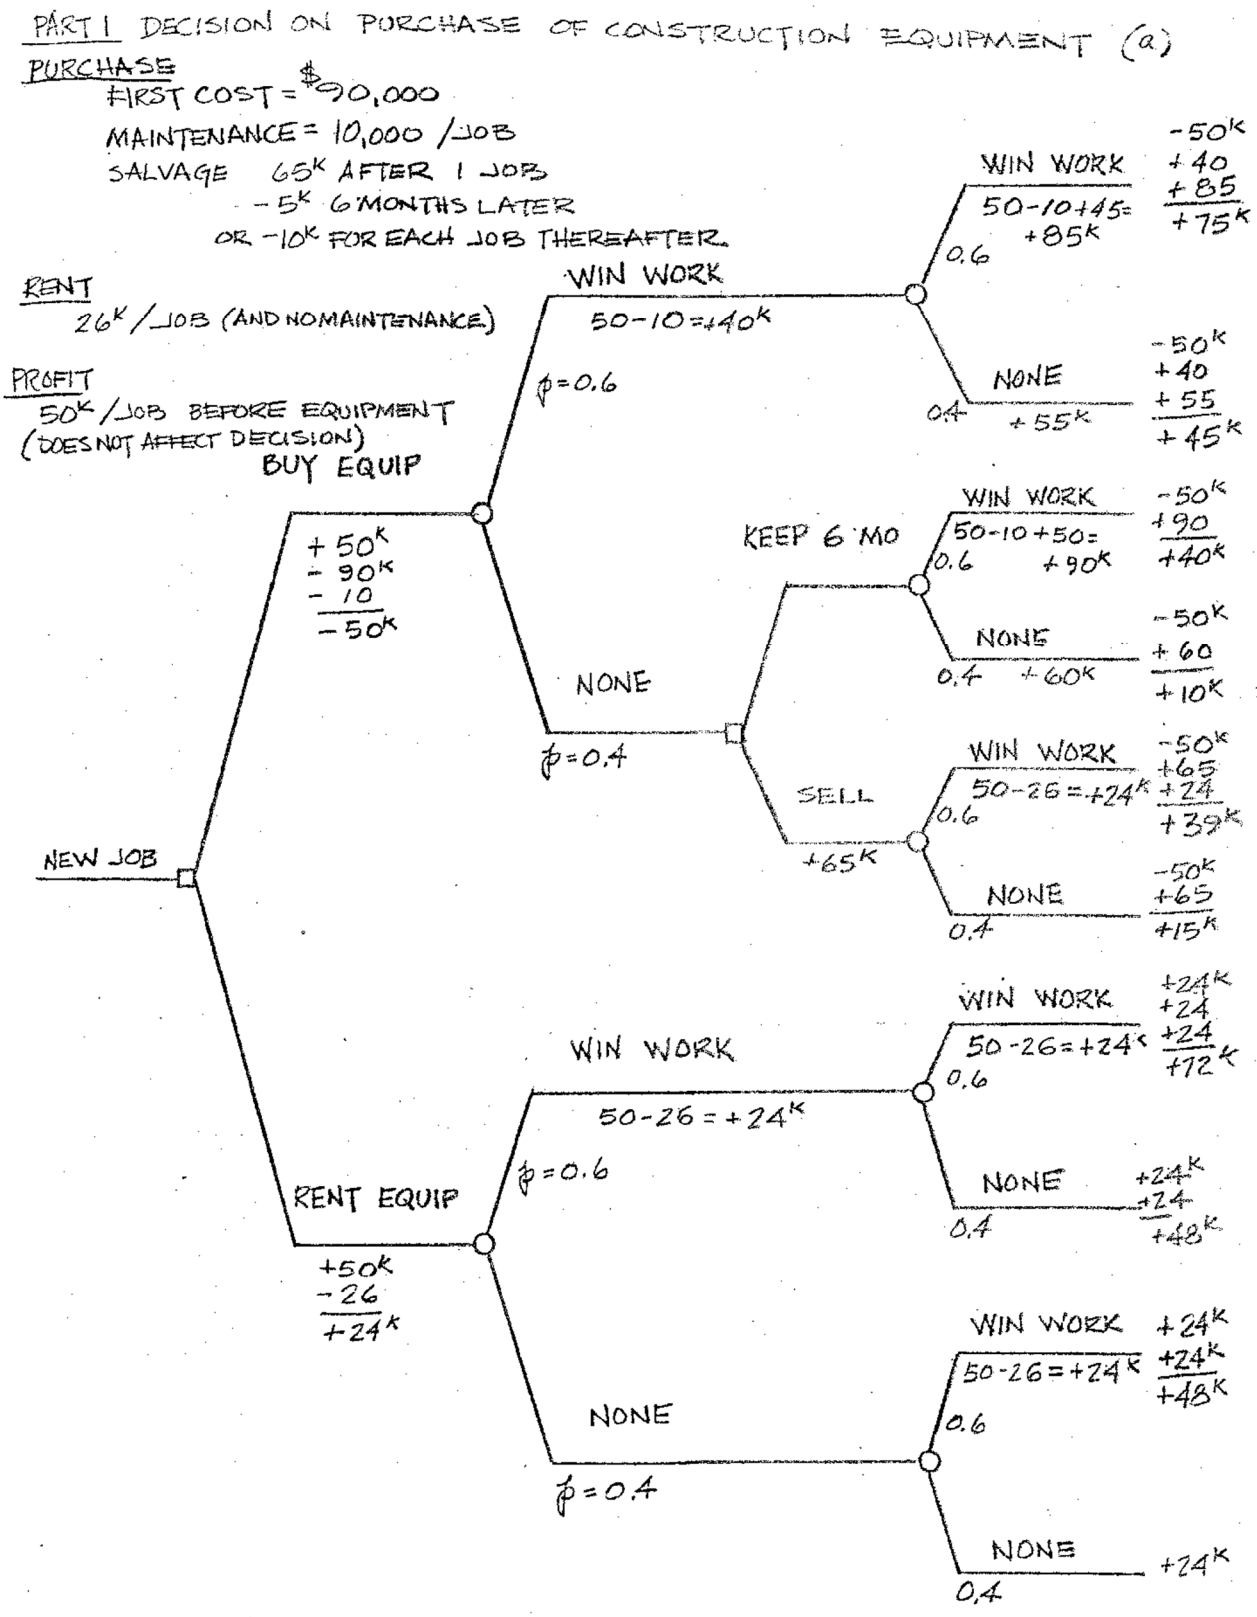


Figure E.8. Simplified Decision to Purchase or Rent Equipment

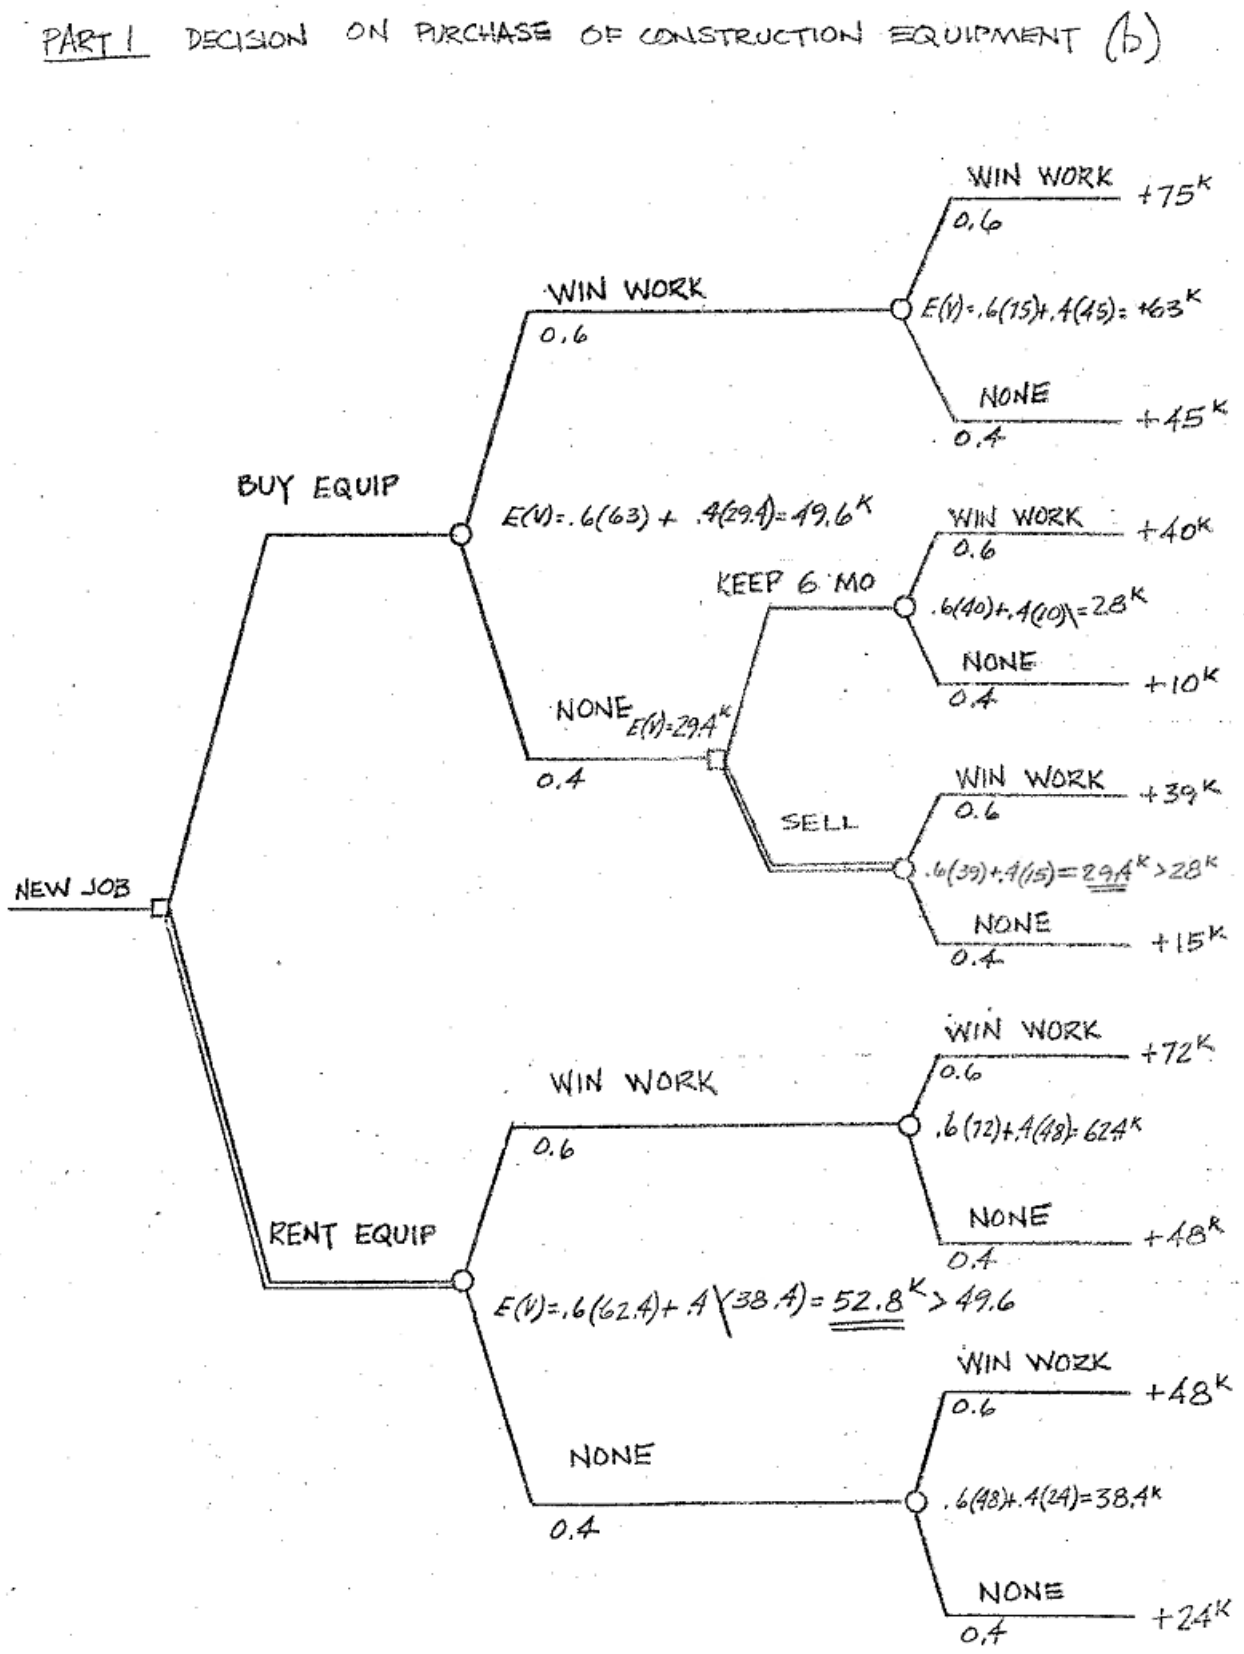


Figure E.9. Basic Decision to Purchase or Rent Equipment

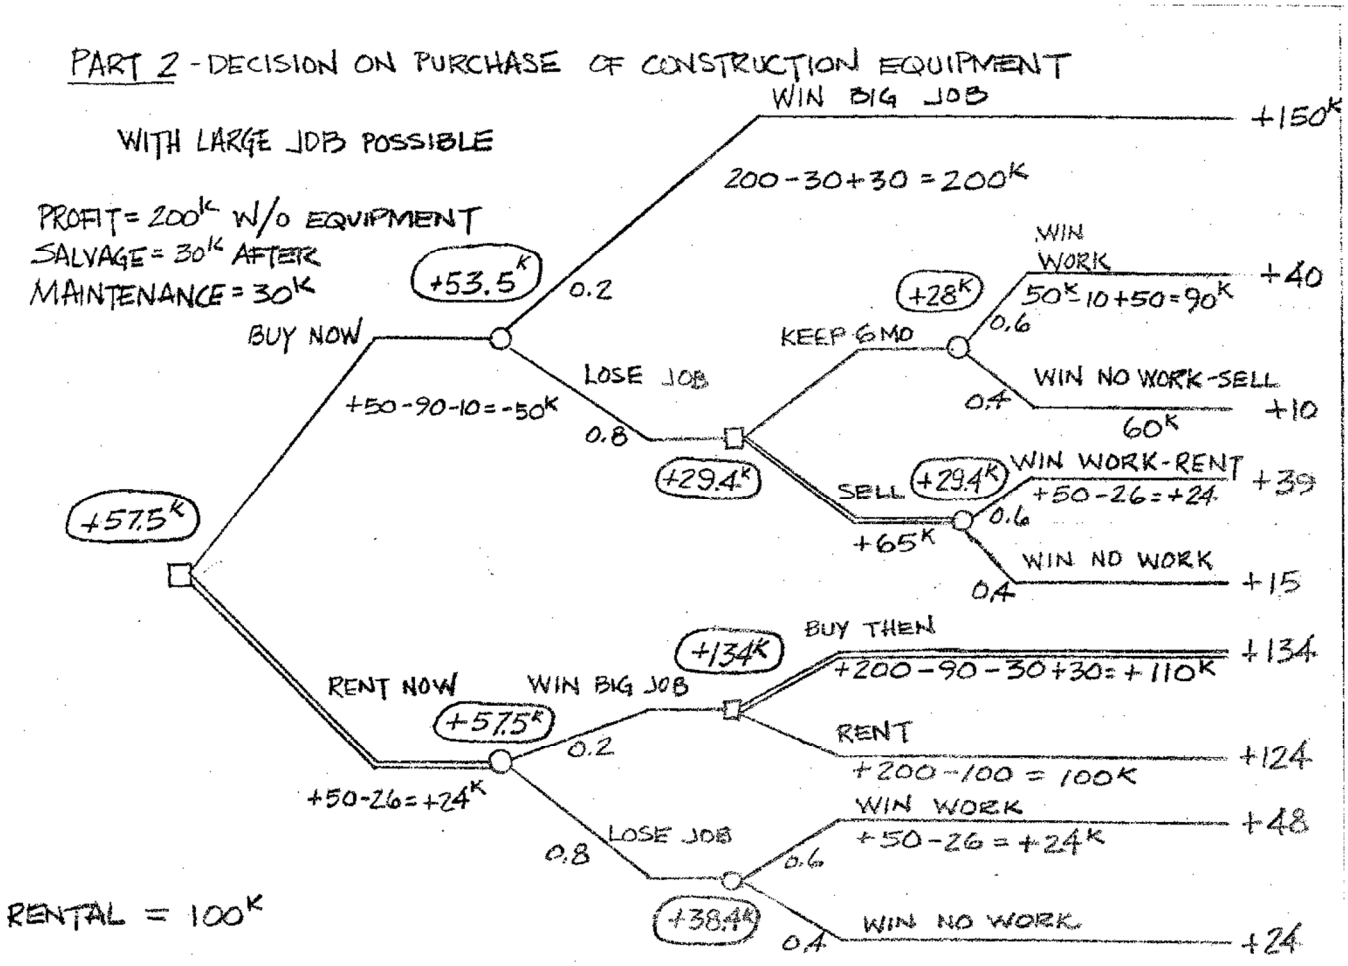


Figure E.10. Decision to Purchase or Rent Equipment with Large Job Option

SIMPLIFIED DECISION TO PURCHASE CONSTRUCTION EQUIPMENT

Equipment Data:
First Cost = $90,000
Maintenance Cost = $10,000 per job
Salvage Value = $65,000 after one job;
-$5,000 6 months later;
-$10,000 for each job thereafter

Profit = $50,000 before equipment cost
Rental Cost = $26,000 per job

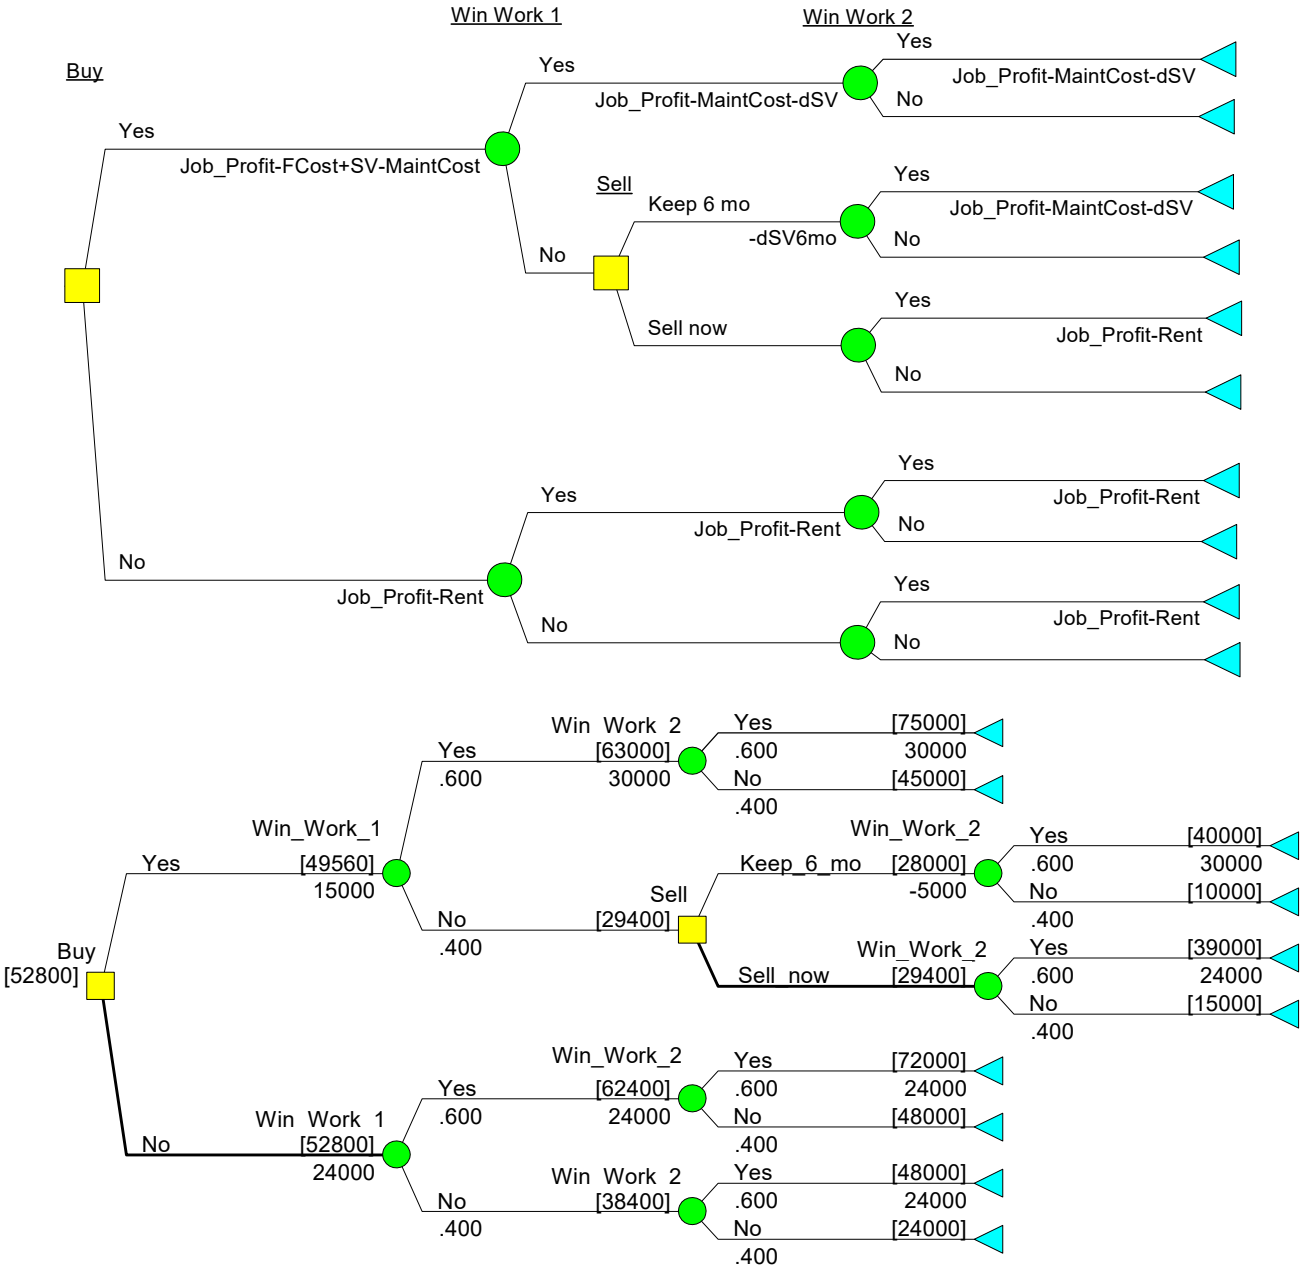


Figure E.11. DPL Model for Basic Decision to Purchase or Rent Equipment

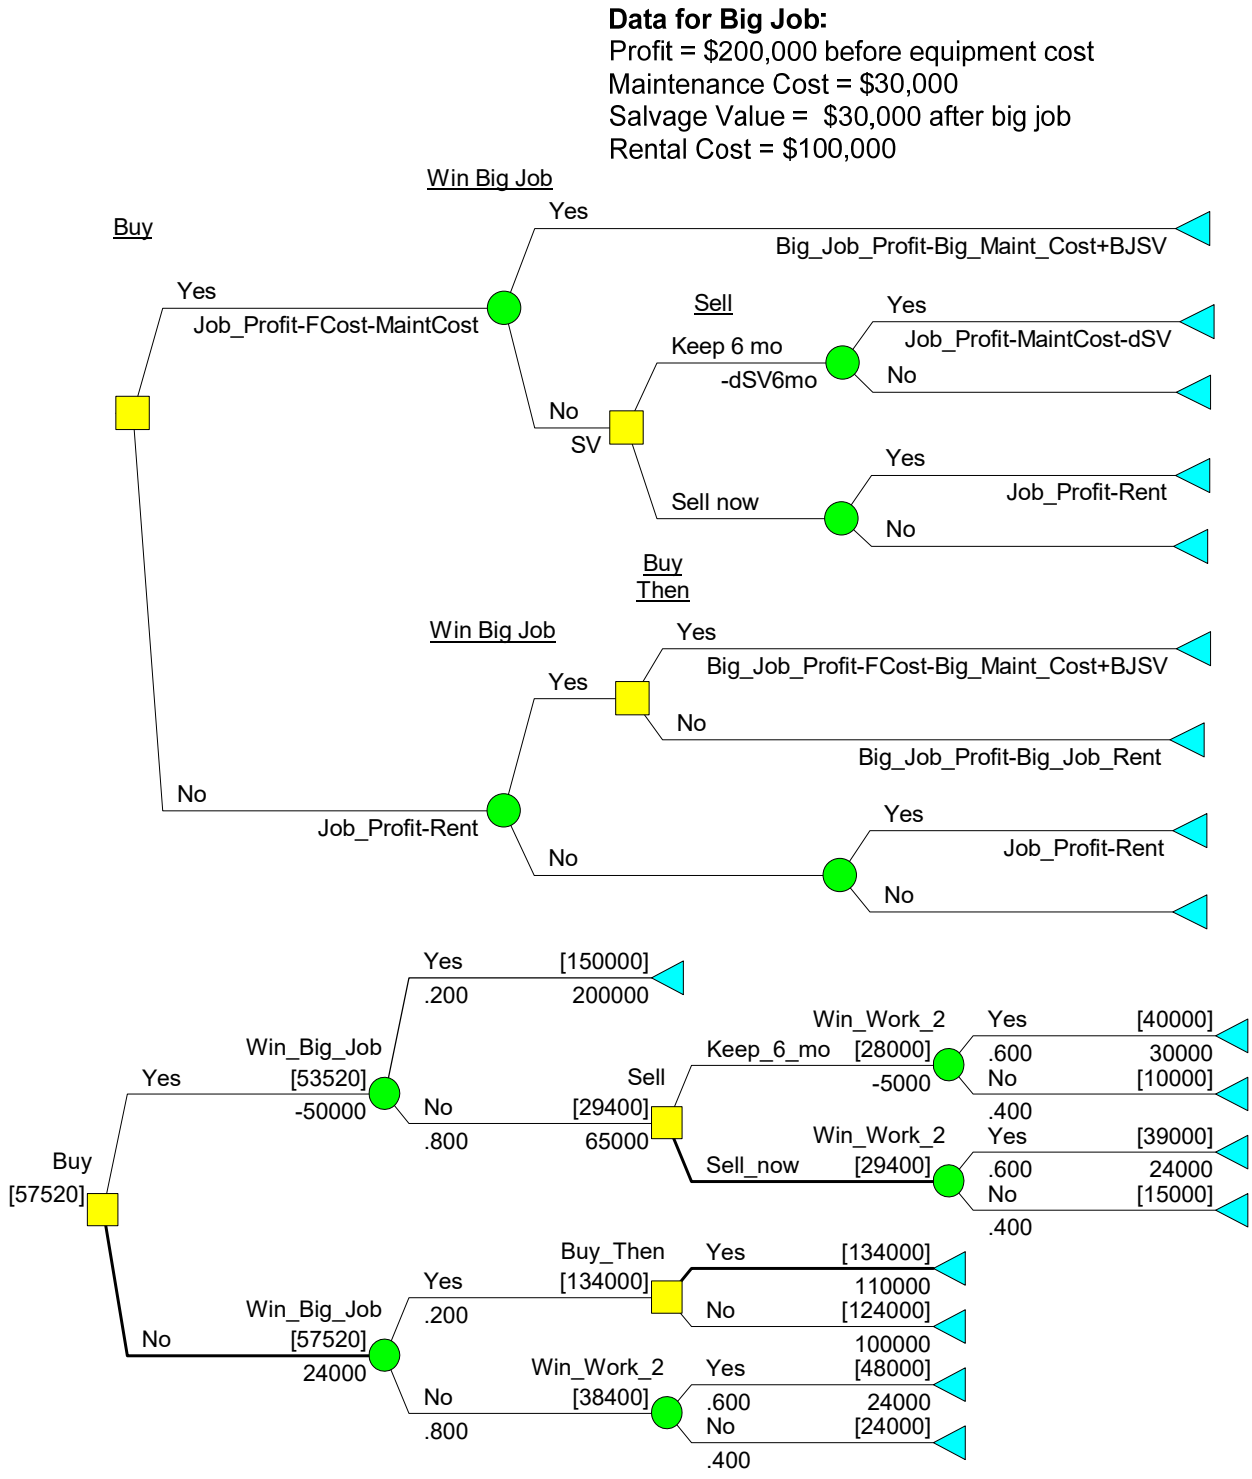


Figure E.12. DPL Model for Decision to Purchase or Rent Equipment with Large Job Option

Important Decision Tree Guidelines

Decision trees must communicate to others the important elements of a decision clearly.

- Complex decision problems may require multiple progressive decision trees to show all details clearly (such as the 6 trees shown in Figures E2 to E7 for the Prefab Plant).
- Each decision node in a decision tree must have a name, such as "Invest", "Test", etc.
- Each chance node must also have a name, such as "Demand in 1st 4 years", etc.
- Each branch at a decision node should have a name, such as "Yes", "No", "Buy", etc.
- Each branch at a chance node should have a name, such as "High", "Low", "S" (Success), "N" (Not Success), "IS" (Indicates Success), "IN" (Indic. Not Success), etc.
- Decision branches must show the associated present values ($): The branch for each decision alternative at a decision node should have a dollar value (in present value dollars) if there are monetary outcomes at that decision branch. This should be a single present dollar value (positive or negative). Any required calculations to convert future cashflows to present value should be done separately and be clearly labeled.
- Chance branches must show the associated present values ($): The branch for each chance outcome at a chance node should have a dollar value (in present value dollars) if there are monetary outcomes at that chance node branch. This should be a single present dollar value (positive or negative). Any required calculations to convert future cashflows to present value should be done separately and be clearly labeled.
- End nodes must show the sum of all associated present values ($): All present values (from all previous decision branches and all previous chance branches) that are on the same path that leads to an end-node should be added up algebraically and appear as a single present dollar value (positive or negative) at the end of the tree to the right of each end node. The end node values are needed to "roll back" the tree (i.e., to calculate the expected monetary values for all preceding nodes). Showing the end node present values clearly is extremely important for decision trees that include risk aversion (a future subject).
- Chance branches must show the associated probabilities: Each chance branch must have an associated probability (single number) next to it. Any required calculations to find the appropriate probability should be done separately and be clearly explained.
- All nodes (branch nodes and decision nodes, including the first node in a decision tree) should show the associated expected monetary values (EMV) on their right. The EMV value at each node should appear on the right side of the node. (Note that the program DPL shows EMV values to the left side of nodes. To differentiate them from branch $ values, i.e., get and pay $ outcomes, DPL encloses EMV values in square brackets.)
- The branch for the best (i.e., the preferred) decision alternative at each decision node should be highlighted (using double lines, red lines, heavy lines, etc.) to show clearly which is the best decision alternative at each decision node. To make the preferred best alternative even clearer, draw two short double lines to "cut off" the branches for the alternatives that are not the best, to show that "we should not go that way".

It is important that all trees follow these guidelines. Please consult these notes for examples.

Sensitivity Analysis

The optimal decision policy for any decision tree depends on the

- Probabilities of future events
- Values of future events

Either of these may be uncertain. We can, of course, estimate the uncertainty and use the estimated probabilities to get expected values (Bayesian statistics). A simpler alternative is to use sensitivity analysis to investigate the changes in optimal decision policy as these parameters change (what-if analysis).

Sensitivity in a decision tree can be shown as

(a) Change in expected value per change in probability or outcome = dy/dx

(b) Magnitude of probability or outcome at which decision will change.

The standard method involves trial calculations using differing magnitudes. For example, at what $P[H_1]$ would the expected value for the large and small plants be equal in Fig. E.7?

$$392.0\ P[H1] – 82.6\ P[L1] = 363.5\ P[H1] + 89.2\ P[L1]$$

substituting $P[L1] = 1 - P[H1]$

$$392.0\ P[H1] -82.6\ (l-P[H1]) = 363.5\ P[H1] - 89.2\ (1-P[H1]) = 0$$

$$(392+82.6-363.5+89.2)\ P[H1] = +82.6 + 89.2$$

$$P[H_1] = 171.8 / 200.3 = 0.86$$

The decision balance between constructing a. large plant and a small plant is at

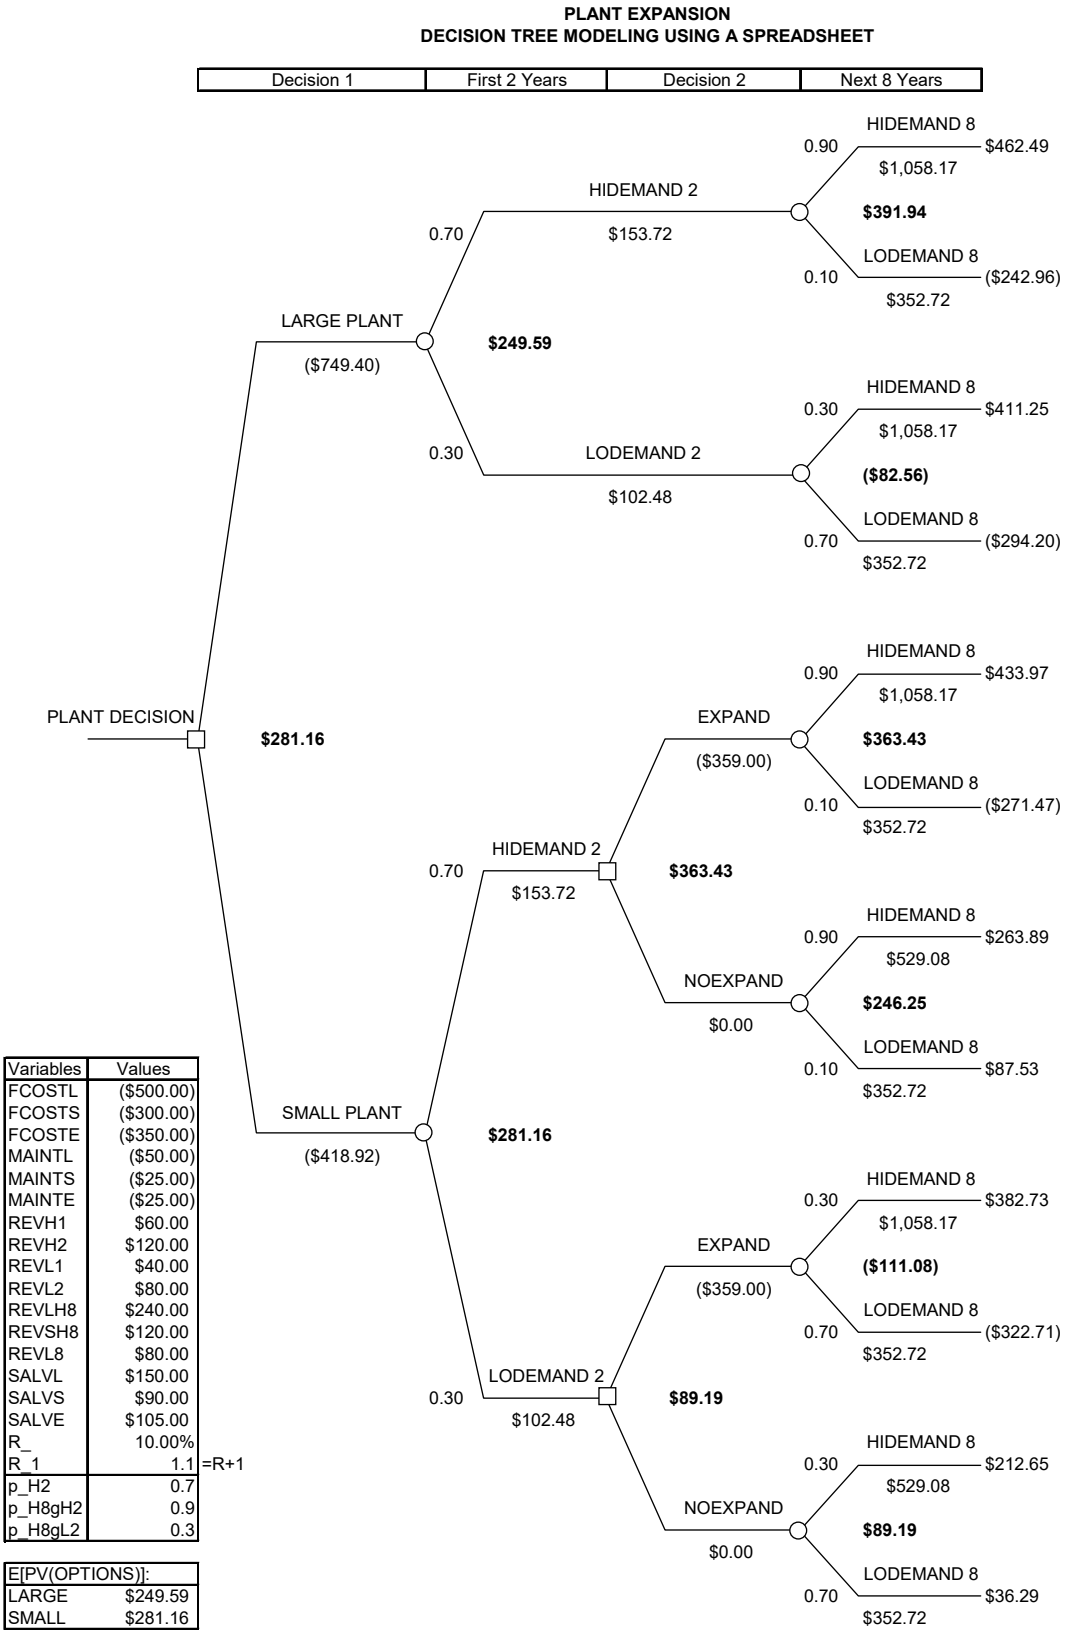
$$P[H_1] = 0.86,\ P[L_1] = 0.14$$

Other possible questions:

- At what labor cost would the decision be indifferent?
- At what purchase price would the decision be indifferent?
- What if the price of lumber goes up? Costs of trusses, whether manufactured by you or purchased will go up. Probably, the decision at hand is not too sensitive to lumber prices.
- What if a competitor cuts his price in response to your competition? The decision at hand may be very sensitive to this possibility.

Sensitivity to rise in wages is less important than sensitivity to uncertainty in your labor production rates (which only affects the cost of your product and not that of your competitors).

Moreover, sensitivity depends on the correlation among random variables, and this may also be investigated.



Variables	Values	
FCOSTL	($500.00)	
FCOSTS	($300.00)	
FCOSTE	($350.00)	
MAINTL	($50.00)	
MAINTS	($25.00)	
MAINTE	($25.00)	
REVH1	$60.00	
REVH2	$120.00	
REVL1	$40.00	
REVL2	$80.00	
REVLH8	$240.00	
REVSH8	$120.00	
REVL8	$80.00	
SALVL	$150.00	
SALVS	$90.00	
SALVE	$105.00	
R_	10.00%	
R_1	1.1	=R+1
p_H2	0.7	
p_H8gH2	0.9	
p_H8gL2	0.3	

E[PV(OPTIONS)]:	
LARGE	$249.59
SMALL	$281.16

Figure E.13. Parametric Decision Tree for Plant Example

PLANT EXPANSION
DECISION TREE MODELING USING A SPREADSHEET
SENSITIVITY ANALYSIS

One-way Sensitivity Analysis

Rate	Large Plant	Small Plant	Best
0%	$873.60	$865.70	L
1%	$785.74	$782.30	L
2%	$705.08	$705.97	S
3%	$630.92	$636.03	S
4%	$562.65	$571.86	S
5%	$499.71	$512.92	S
6%	$441.63	$458.71	S
7%	$387.95	$408.81	S
8%	$338.28	$362.81	S
9%	$292.27	$320.37	S
10%	$249.59	$281.16	S

Two-way Sensitivity Analysis

Rate	p_H2										
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0%	S	S	S	S	S	S	S	L	L	L	L
1%	S	S	S	S	S	S	S	L	L	L	L
2%	S	S	S	S	S	S	S	S	L	L	L
3%	S	S	S	S	S	S	S	S	L	L	L
4%	S	S	S	S	S	S	S	S	L	L	L
5%	S	S	S	S	S	S	S	S	L	L	L
6%	S	S	S	S	S	S	S	S	L	L	L
7%	S	S	S	S	S	S	S	S	S	L	L
8%	S	S	S	S	S	S	S	S	S	L	L
9%	S	S	S	S	S	S	S	S	S	L	L
10%	S	S	S	S	S	S	S	S	S	L	L

p_H2	Large Plant	Small Plant	Best
0	($82.56)	$89.19	S
0.1	($35.11)	$116.62	S
0.2	$12.34	$144.04	S
0.3	$59.79	$171.46	S
0.4	$107.24	$198.89	S
0.5	$154.69	$226.31	S
0.6	$202.14	$253.73	S
0.7	$249.59	$281.16	S
0.8	$297.04	$308.58	S
0.9	$344.49	$336.00	L
1	$391.94	$363.43	L

p_H2	p_H8gL2										
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	S	S	S	S	S	S	S	L	L	L	L
0.1	S	S	S	S	S	S	S	L	L	L	L
0.2	S	S	S	S	S	S	S	L	L	L	L
0.3	S	S	S	S	S	S	S	L	L	L	L
0.4	S	S	S	S	S	S	L	L	L	L	L
0.5	S	S	S	S	S	S	L	L	L	L	L
0.6	S	S	S	S	S	S	L	L	L	L	L
0.7	S	S	S	S	S	L	L	L	L	L	L
0.8	S	S	S	S	S	L	L	L	L	L	L
0.9	S	S	L	L	L	L	L	L	L	L	L
1	L	L	L	L	L	L	L	L	L	L	L

p_H8gH2	Large Plant	Small Plant	Best
0	($194.84)	$88.03	S
0.1	($145.46)	$100.37	S
0.2	($96.08)	$112.72	S
0.3	($46.69)	$125.06	S
0.4	$2.69	$137.41	S
0.5	$52.07	$149.75	S
0.6	$101.45	$162.10	S
0.7	$150.83	$182.39	S
0.8	$200.21	$231.78	S
0.9	$249.59	$281.16	S
1	$298.97	$330.54	S

p_H8gH2	p_H2										
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	S	S	S	S	S	S	S	S	S	S	S
0.1	S	S	S	S	S	S	S	S	S	S	S
0.2	S	S	S	S	S	S	S	S	S	S	S
0.3	S	S	S	S	S	S	S	S	S	S	S
0.4	S	S	S	S	S	S	S	S	S	S	S
0.5	S	S	S	S	S	S	S	S	S	S	S
0.6	S	S	S	S	S	S	S	S	S	S	S
0.7	S	S	S	S	S	S	S	S	S	L	L
0.8	S	S	S	S	S	S	S	S	S	L	L
0.9	S	S	S	S	S	S	S	S	S	L	L
1	S	S	S	S	S	S	S	S	S	L	L

p_H8gL2	Large Plant	Small Plant	Best
0	$186.10	$265.28	S
0.1	$207.27	$270.58	S
0.2	$228.43	$275.87	S
0.3	$249.59	$281.16	S
0.4	$270.76	$286.45	S
0.5	$291.92	$291.74	L
0.6	$313.08	$297.03	L
0.7	$334.25	$305.73	L
0.8	$355.41	$326.89	L
0.9	$376.57	$348.05	L
1	$397.74	$369.22	L

p_H8gL2	p_H2										
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	S	S	S	S	S	S	S	S	S	S	L
0.1	S	S	S	S	S	S	S	S	S	S	L
0.2	S	S	S	S	S	S	S	S	S	L	L
0.3	S	S	S	S	S	S	S	S	S	L	L
0.4	S	S	S	S	S	S	S	S	S	L	L
0.5	S	S	S	S	S	S	S	L	L	L	L
0.6	S	S	S	S	L	L	L	L	L	L	L
0.7	L	L	L	L	L	L	L	L	L	L	L
0.8	L	L	L	L	L	L	L	L	L	L	L
0.9	L	L	L	L	L	L	L	L	L	L	L
1	L	L	L	L	L	L	L	L	L	L	L

Figure E.14. Sensitivity Analysis for Plant Example Using Excel Tables

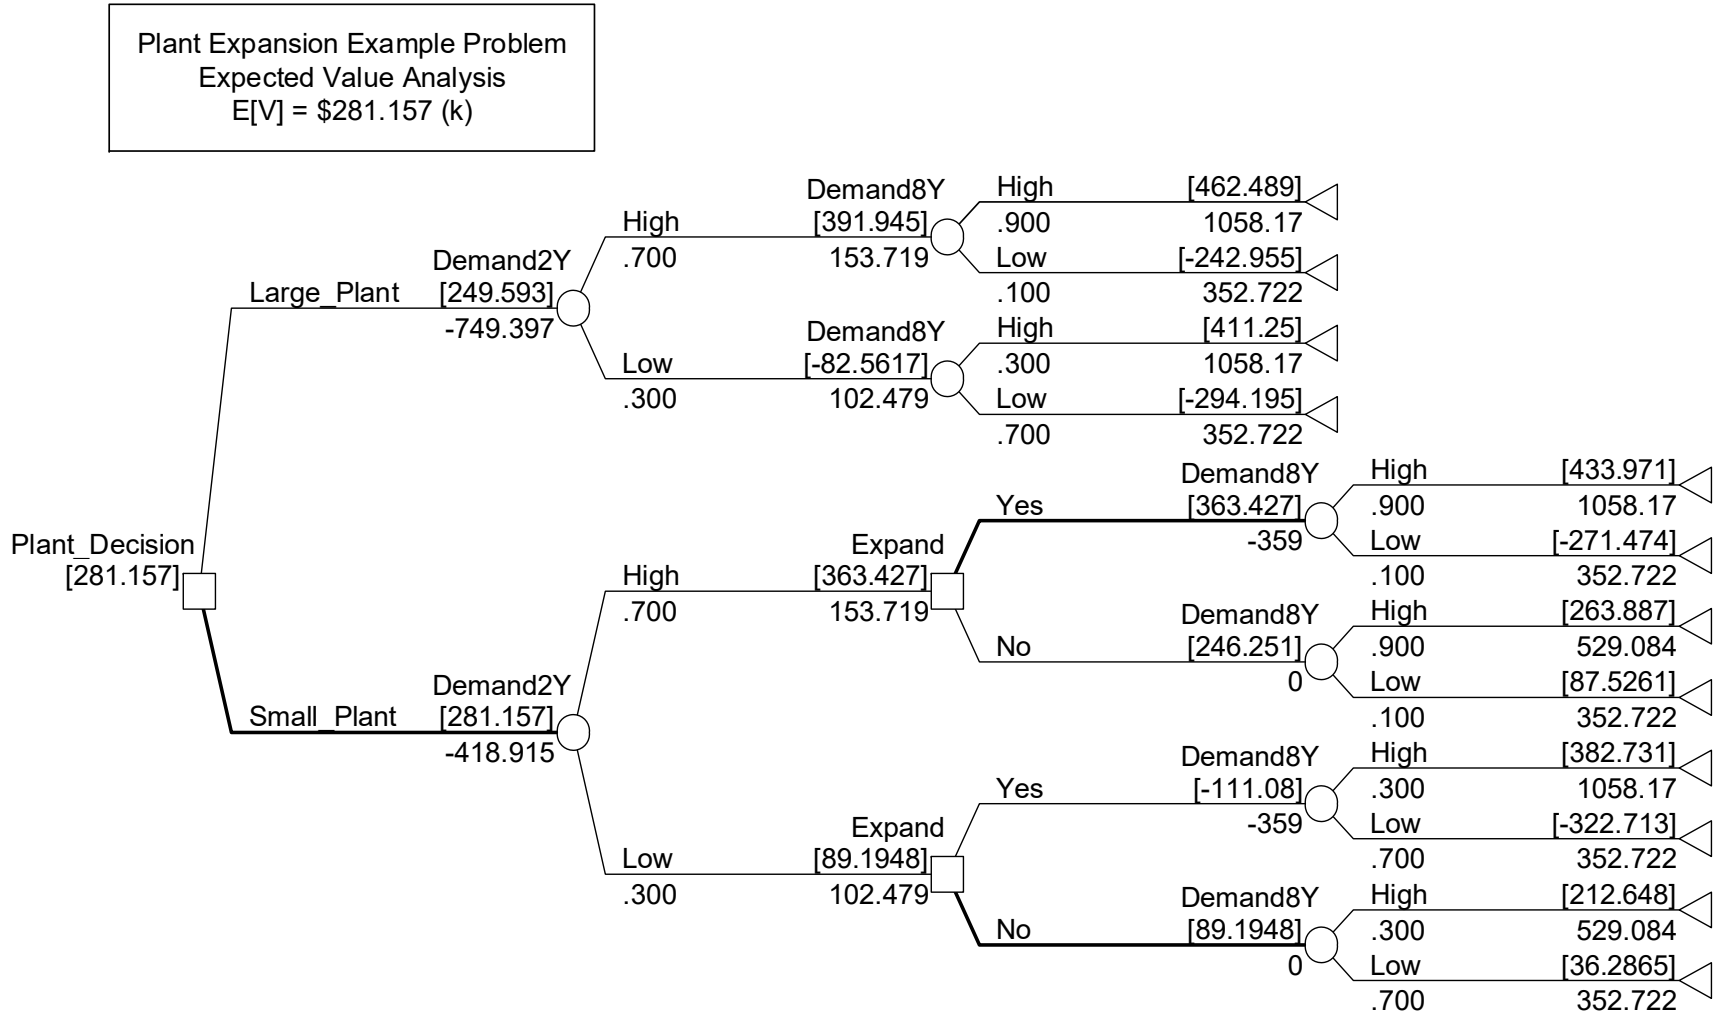


Figure E.15. DPL Model for Plant Example

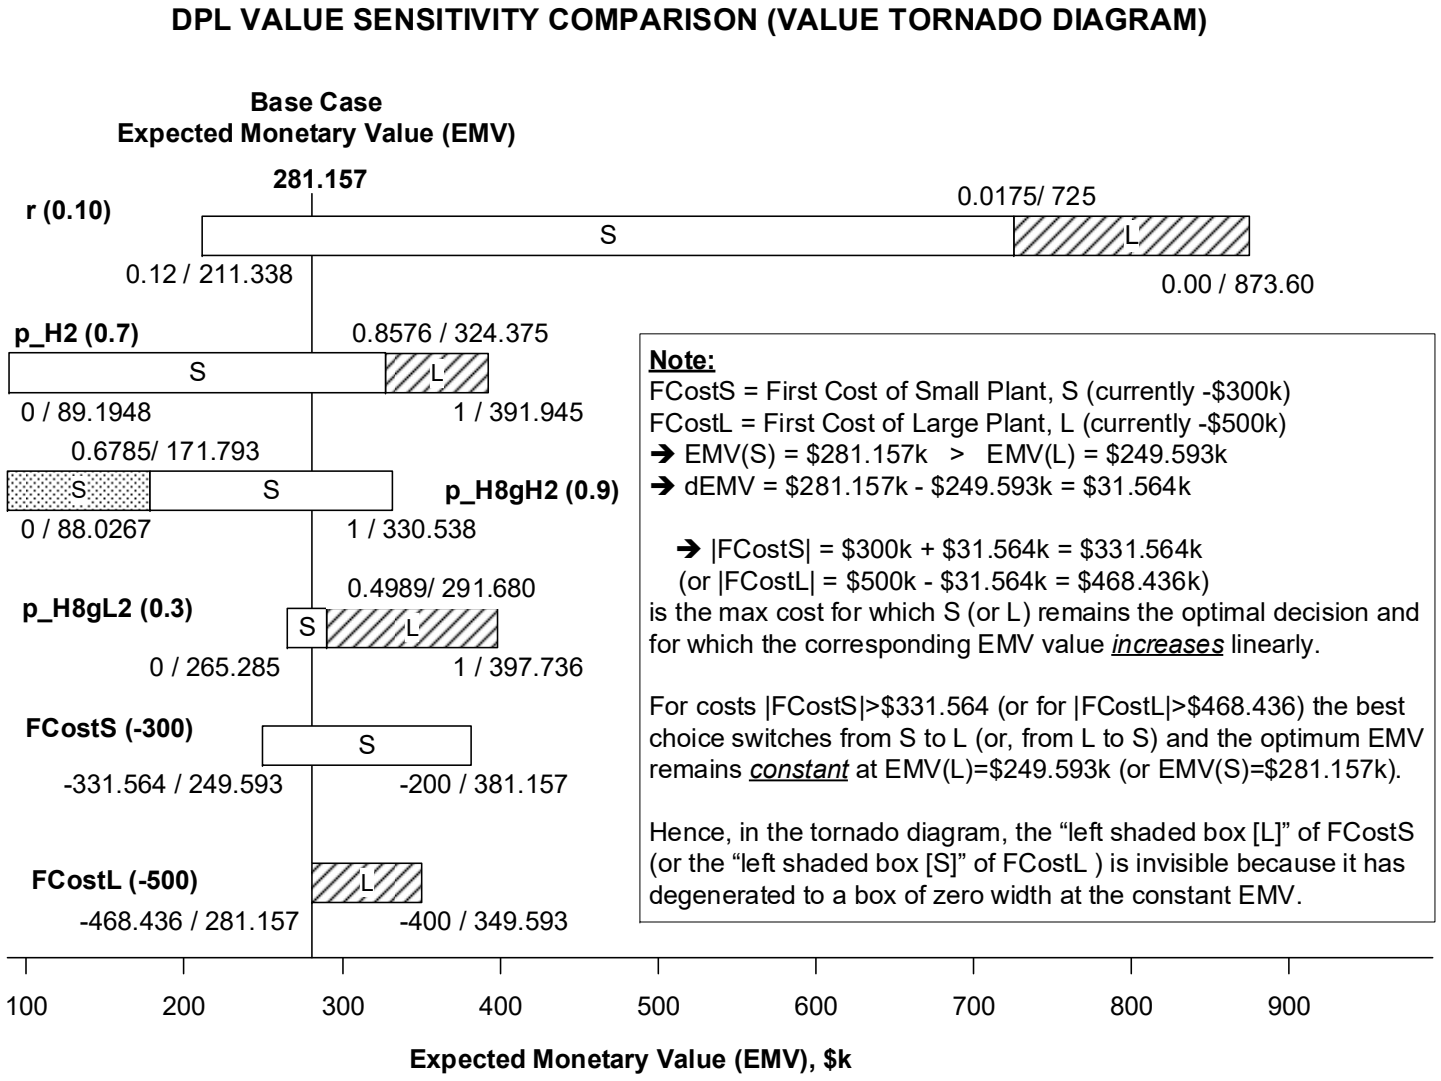


Figure E.16. Sensitivity Analysis for Plant Example using DPL Tornado Diagram

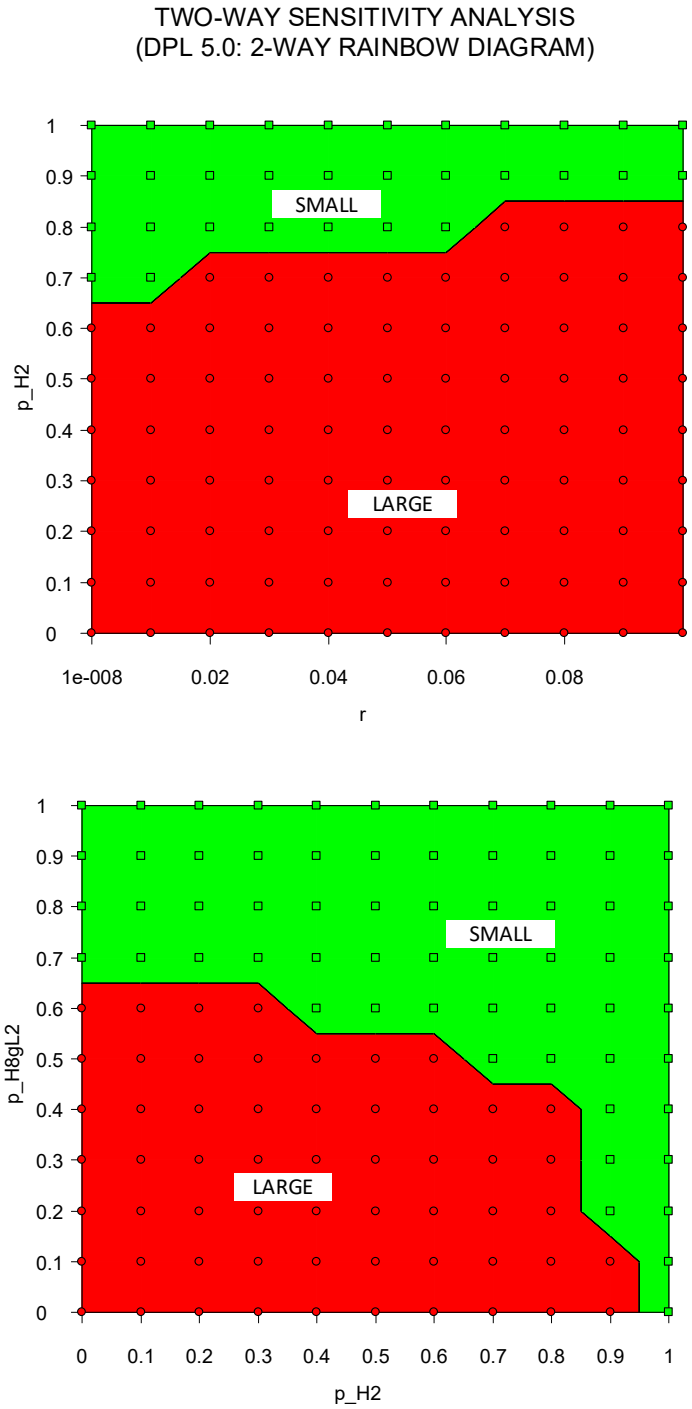


Figure E.17. Two-way Sensitivity Analysis for Plant Example using DPL Rainbow Diagram

Expected Value of Perfect Information (EVPI)

Certainty is knowledge of outcomes with probability one. A decision under certainty is simply a choice between certain outcomes. We know from the study of expected value and decision trees that a preference among alternatives depends upon future outcomes. If the outcomes were known, the proper alternative could be chosen with certainty. Research which predicted outcomes perfectly, delivering perfect information, would therefore be of considerable value in many decisions.

In certain cases, perfect information is only available as hindsight, after what was the future becomes the past. In other cases, such as geologic exploration or the estimation of structural loads or strength, perfect information can be obtained. In either case, calculation of the value of perfect information can be helpful, because it gives an upper bound to the value of any research which promises useful information on the outcome of an uncertain variable. Imperfect information is obviously of no greater value than perfect information.

The relative positions of nodes on a decision tree represent the relative times of knowledge of event occurrences rather than the times when the events actually occur. Therefore, a decision made before the outcome of a chance node is known is shown as a decision node preceding the chance node. Likewise, a decision made after the outcome is known is shown after the chance node, even though the outcome has not occurred prior to the decision.

Decisions made with perfect information of chance events are made after those events are known (but have not yet necessarily occurred). Decision nodes on the tree are placed to the right of those chance nodes.

The calculation of expected value for the perfect information decision tree requires probabilities for the outcomes of the research which produces the perfect information. Since there is perfect correlation between the occurrence of an event and its prediction, the probabilities assigned for the occurrence of the events are also the probabilities of outcome of the research. That is, if the probability of bedrock at 20 ft. is estimated to be 0.3 and a test will predict the actual location with certainty, then there is a 0.3 probability that the test will predict bedrock at 20 ft. The probabilities of the chance events are therefore placed on the outcomes of the perfect research on the tree.

Expected value calculations on the tree produce the expected value of the investment opportunity if perfect information were available on chance events *before* decisions were made. The difference between the expected monetary values, EMV, with and without perfect information is called the *expected value of perfect information* (EVPI). Research which produces less than perfect information will have an expected value equal to or less than that of perfect information. The value of the information provided by an *imperfect* test program is called the *expected value of sampled information* (EVSI). Obviously, EVPI >= EVSI >= $0. Clearly, imperfect research whose *cost* >= EVPI >= EVSI can be dismissed as too expensive without any need to waste effort to determine its EVSI. However, imperfect research might be worth performing if its cost <= EVPI. In this case, we need to perform further analysis to compute its EVSI and then undertake the research only if its cost <= EVSI.

Example. Fig. E.7 shows that building a small plant is preferred to a large plant because it has a higher expected value. However, if a high demand were known to be the outcome, the large plant would be preferred.

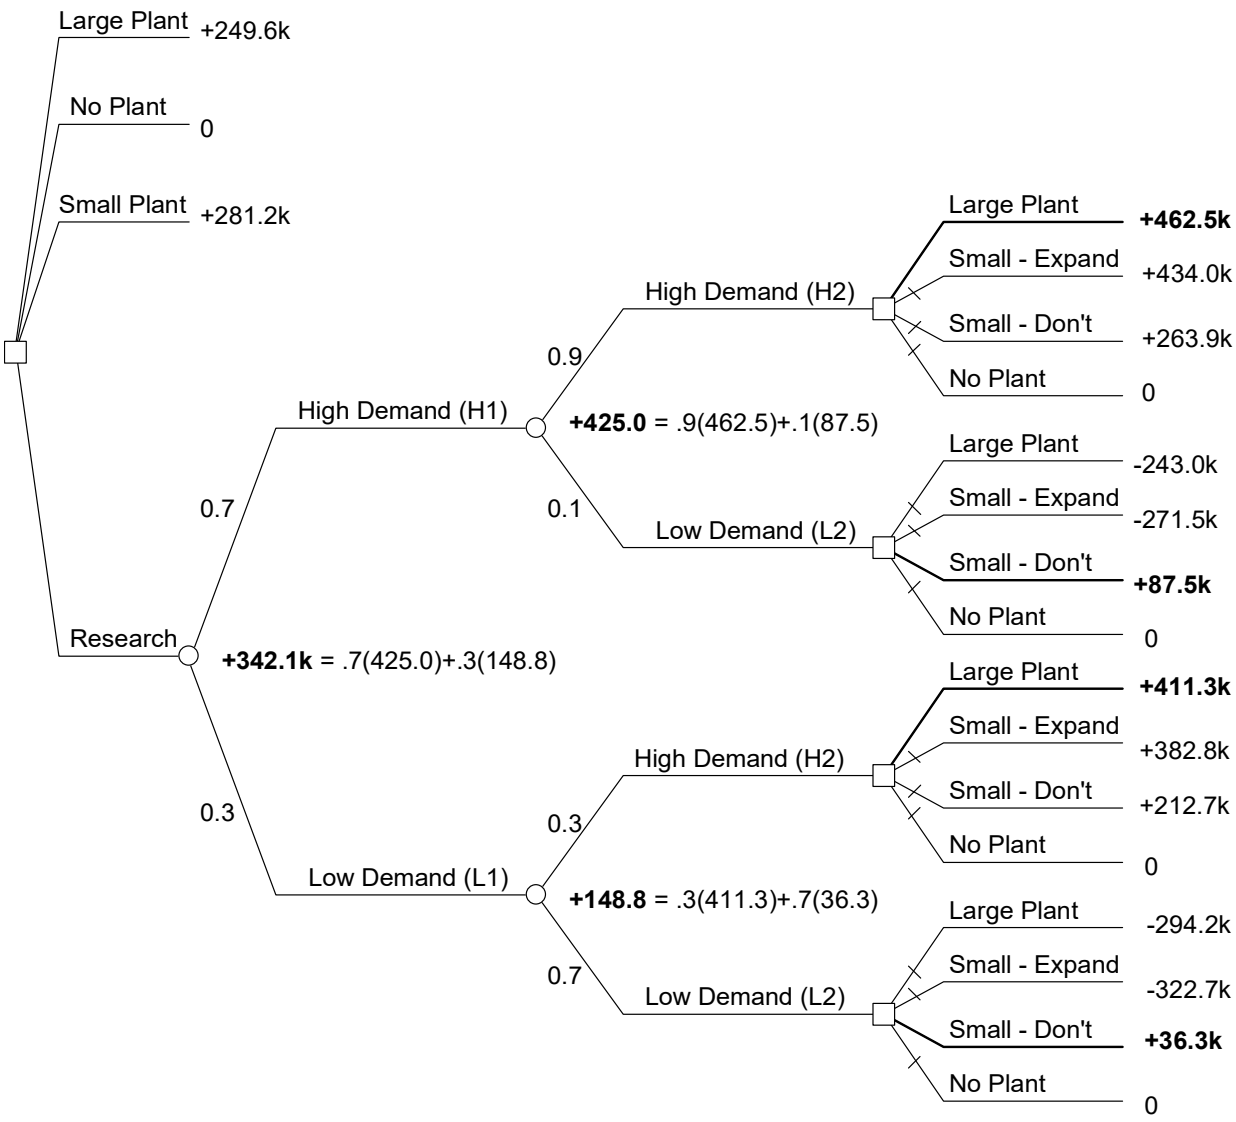
This is easily demonstrated on the decision tree for perfect information shown in Fig. E.18 *where demand for all 10 years will be known before the decision is required*. The chance nodes are shown *prior* to decision of plant size because *the outcomes of those nodes will be known before the decisions are made*. The end-node present values associated with each decision as shown at the far right of the tree are taken from the respective combination of decision and chance events in Fig. E.5. In other words, the end-node present value outcomes in Fig. E.5 and Fig. E.18 are exactly the same, except for the fact that they appear in a different order from top to bottom to match the chance events and decision choices that precede each end-node value.

The probabilities for outcomes of the research in Fig. E.18 are the same as those on the nature's tree shown in Fig. E.6 since there is perfect correlation between occurrence of an event and its prediction. Expected value calculations are shown for the tree. The alternative with the highest value is first selected for each decision node at the right side. Thus, if demand in the first 2 years will be low and demand in the following 8 years high, the large plant should be built with an expected value of $411.3k.

After an expected-value has been determined in this manner for each decision node, the expected value at each earlier node is calculated. This shows that the expected value at the far-left node is $342.1k. This is the expected value when having perfect information. Because this is higher than the best alternative based on present information of $281,200, the value of the perfect information would be

$$\text{EVPI(Demand for 10 years)} = \$342{,}100\text{-}281{,}200 = \$\ 60,\ 900$$

One would spend no more than $60,900 on research to predict the demand over the next 10 years regardless of its accuracy. One might notice that if perfect information were available, the alternative of small plant and later expansion would never be selected although it is the most optimal decision policy without perfect information.



Expected Value of Perfect Information (EVPI)

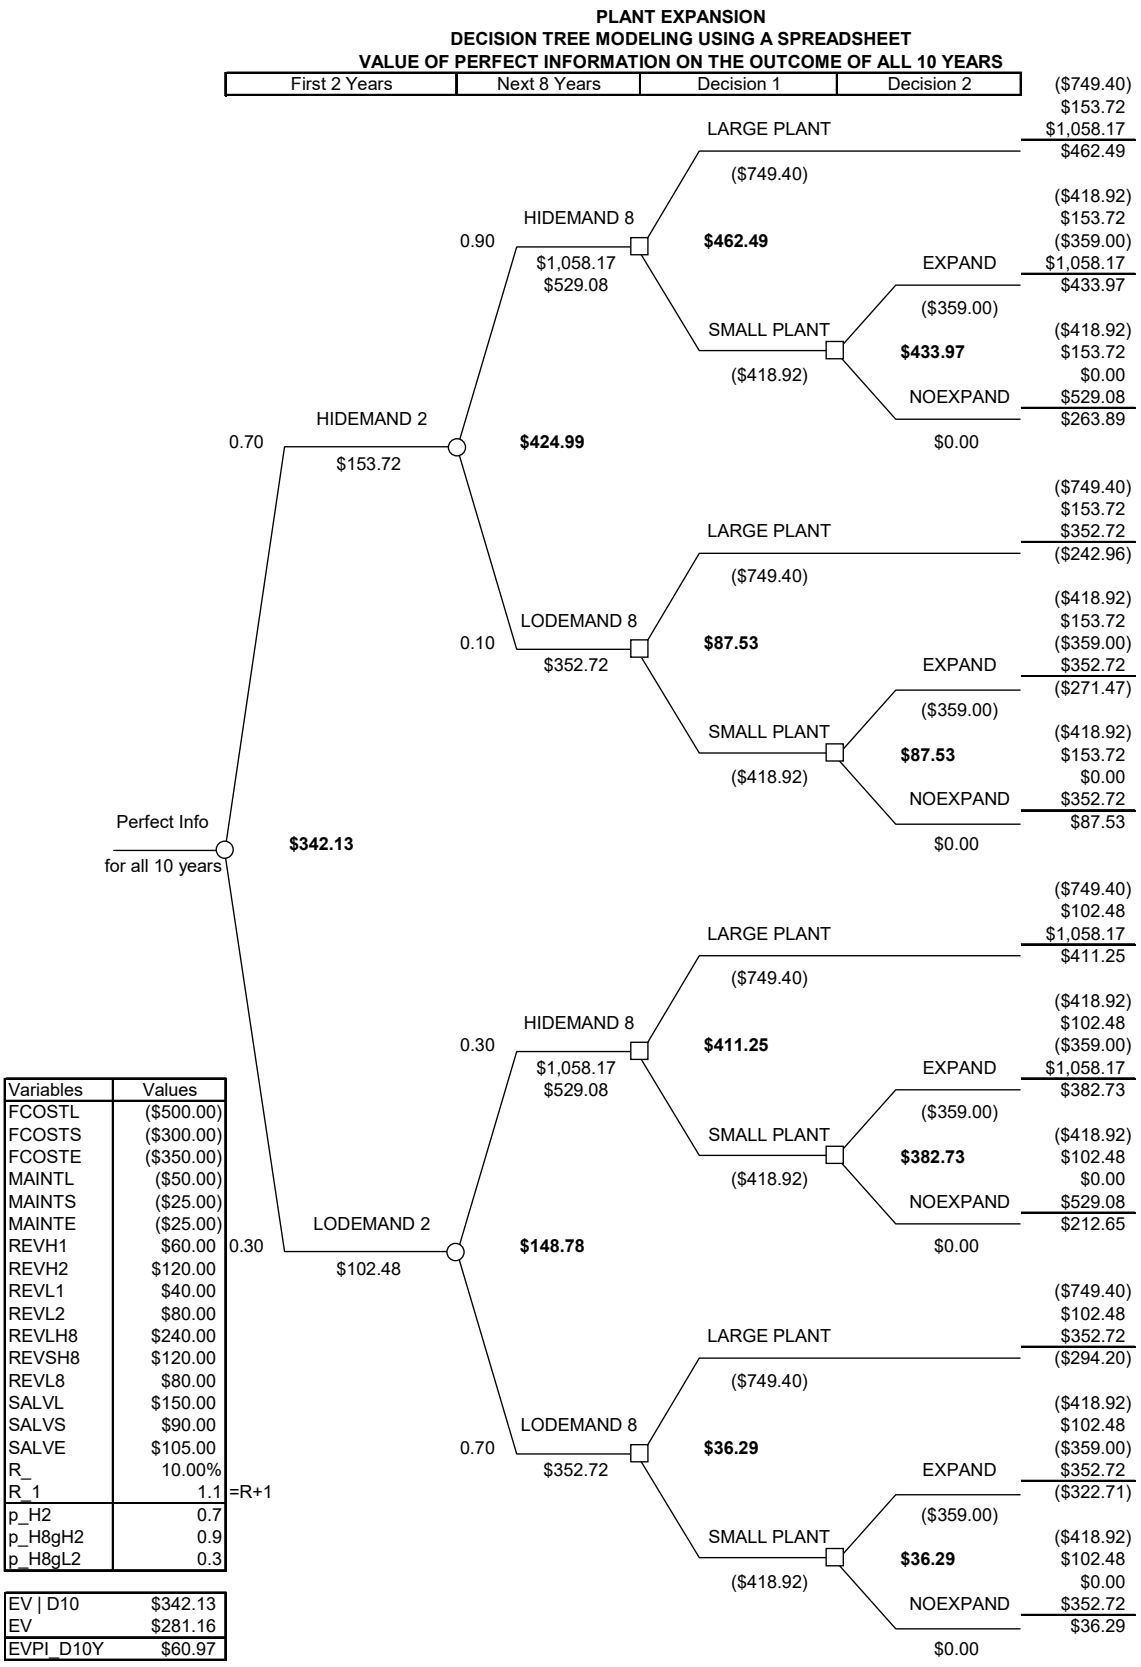
EVPI = Expected Value given Perfect Information - Max Expected Value without Perfect Information

For the case above:
EVPI = 342.1k - 281.2k = $60.9k

In general, perfect research should be undertaken only if its cost, CR, is less than the EVPI. Notice that the cost of the research, CR, is a function of the current state of research technology and is thus determined by the open market. CR is not determined by the decision problem at hand. In contrast, the EVPI is a function of its use in a particular decision problem and thus depends on the state of information of the decision maker.

The EVPI is an upper limit on the value of the information provided by any imperfect research program and thus it is also an upper limit on the cost of such a research program.

Figure E.18. EVPI about Demand Over Ten Years



Variables	Values	
FCOSTL	($500.00)	
FCOSTS	($300.00)	
FCOSTE	($350.00)	
MAINTL	($50.00)	
MAINTS	($25.00)	
MAINTE	($25.00)	
REVH1	$60.00	
REVH2	$120.00	
REVL1	$40.00	
REVL2	$80.00	
REVLH8	$240.00	
REVSH8	$120.00	
REVL8	$80.00	
SALVL	$150.00	
SALVS	$90.00	
SALVE	$105.00	
R_	10.00%	
R_1	1.1	=R+1
p_H2	0.7	
p_H8gH2	0.9	
p_H8gL2	0.3	

EV | D10	$342.13
EV	$281.16
EVPI_D10Y	$60.97

Figure E.19. EVPI about Demand Over Ten Years (Excel)

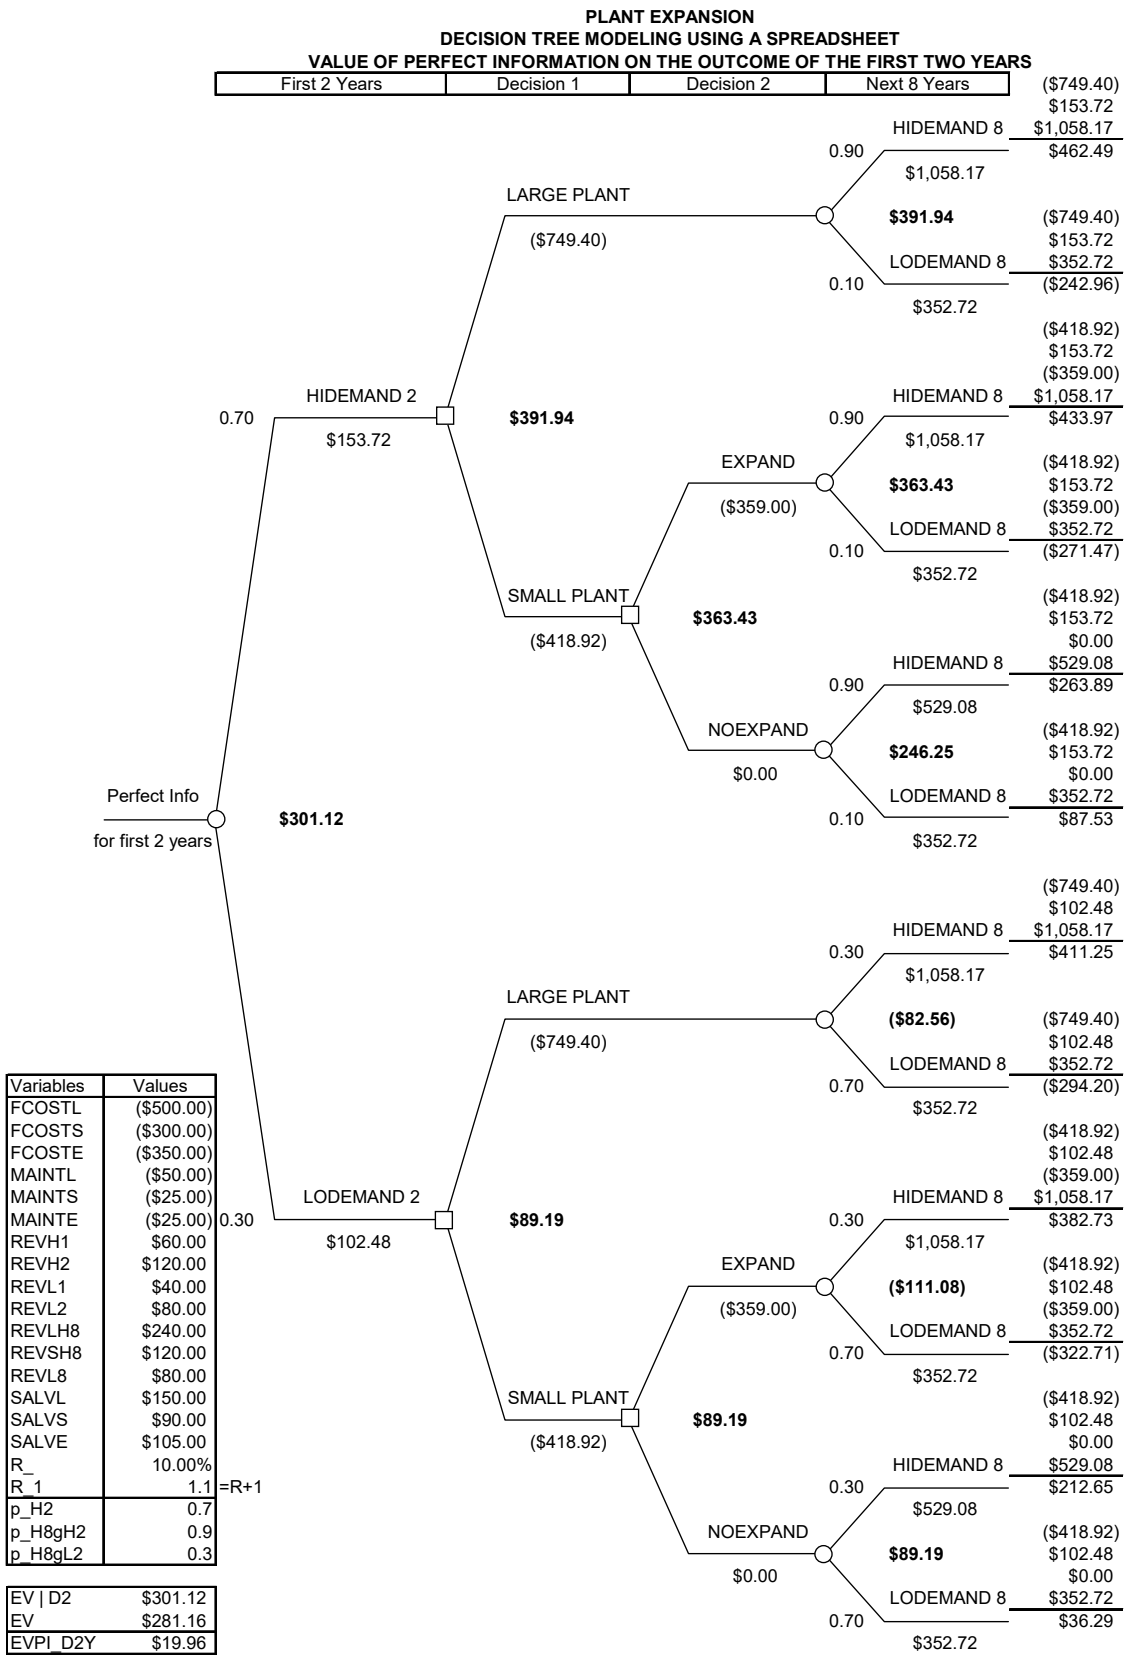


Figure E.20. EVPI about Demand Over First Two Years (Excel)

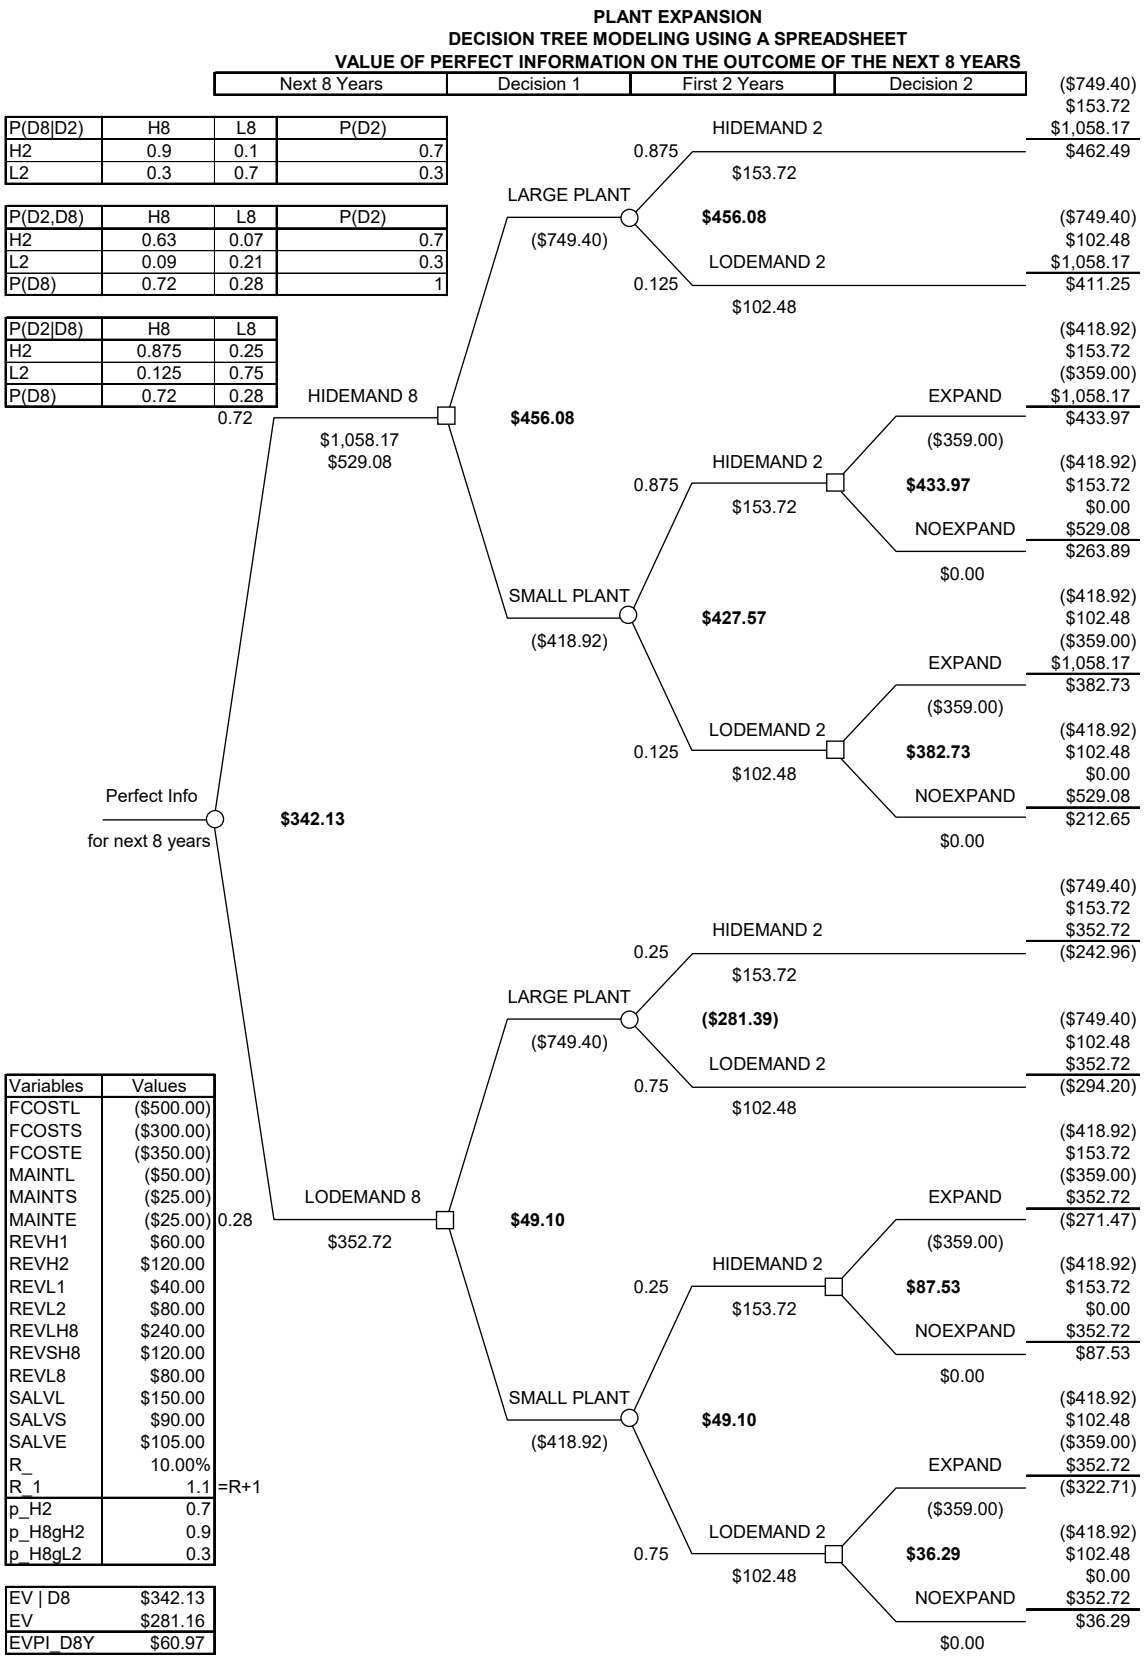


Figure E.21. EVPI about Demand Over Next Eight Years (Excel)

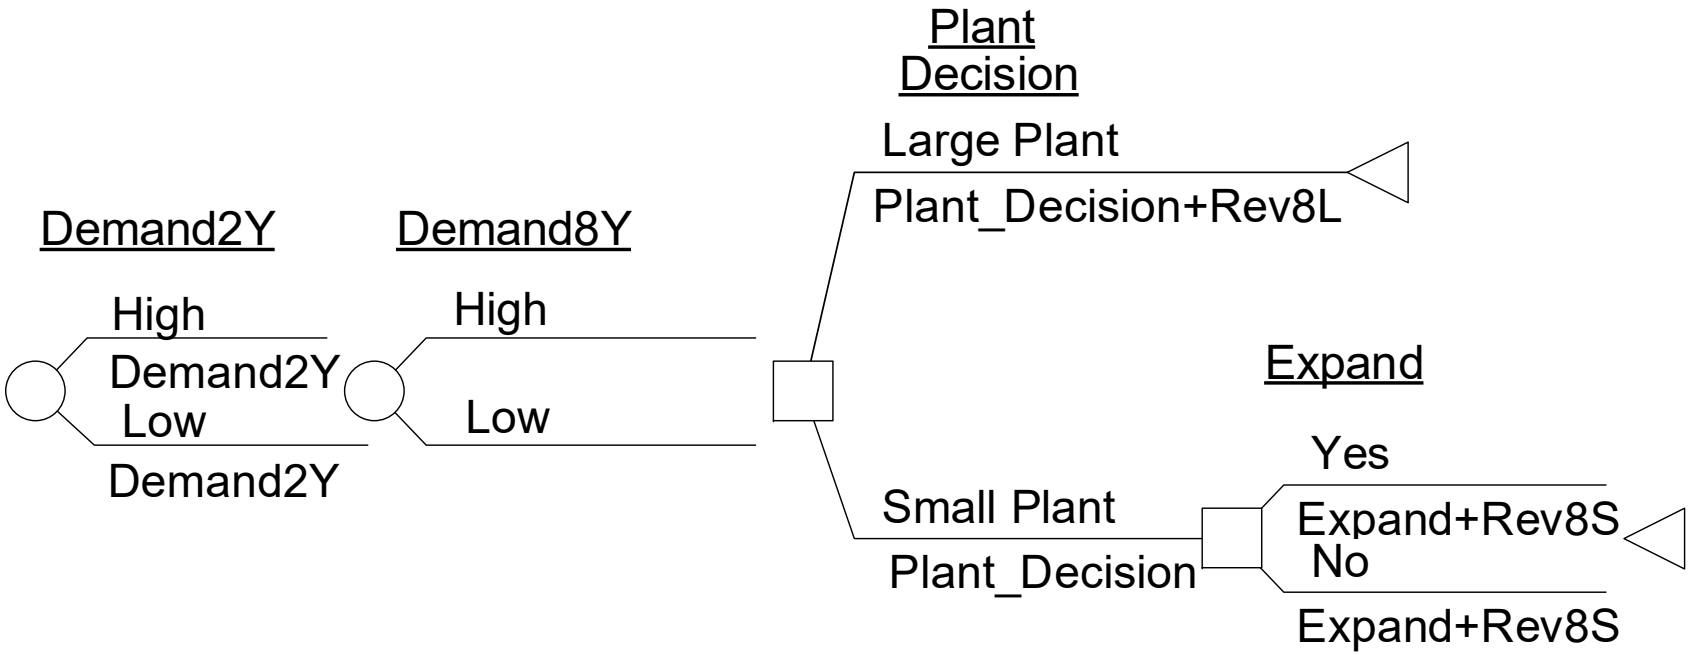


Figure E.22. Generic DPL Decision Tree for EVPI about Demand Over Ten Years

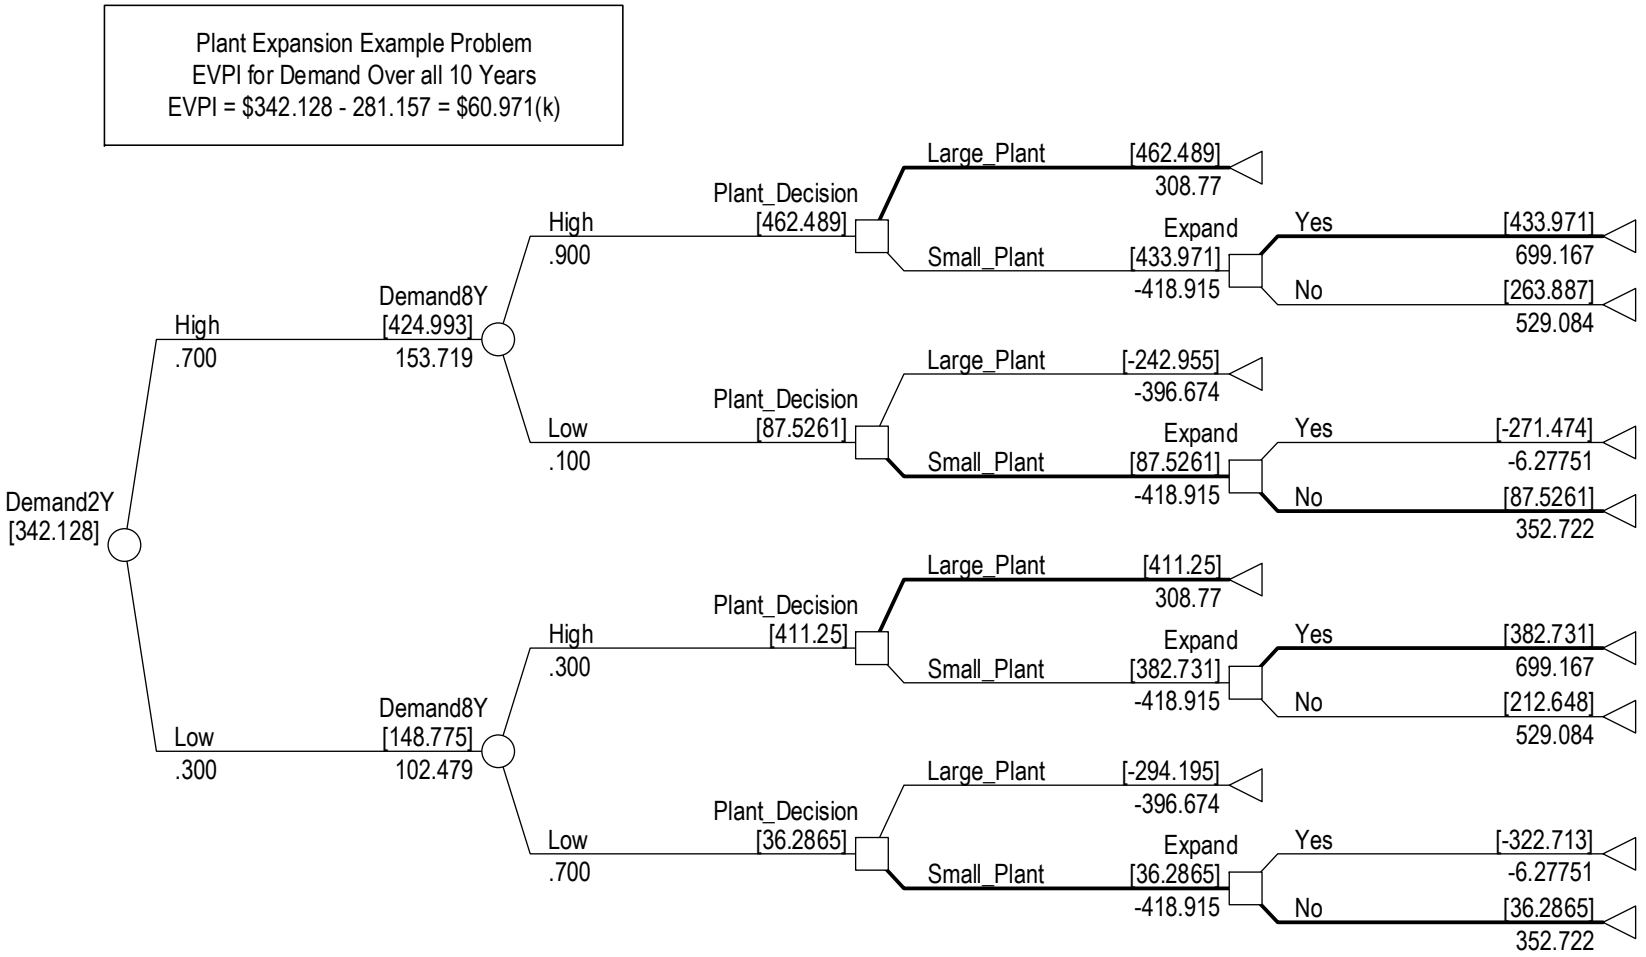


Figure E.23. EVPI about Demand Over Ten Years (DPL)

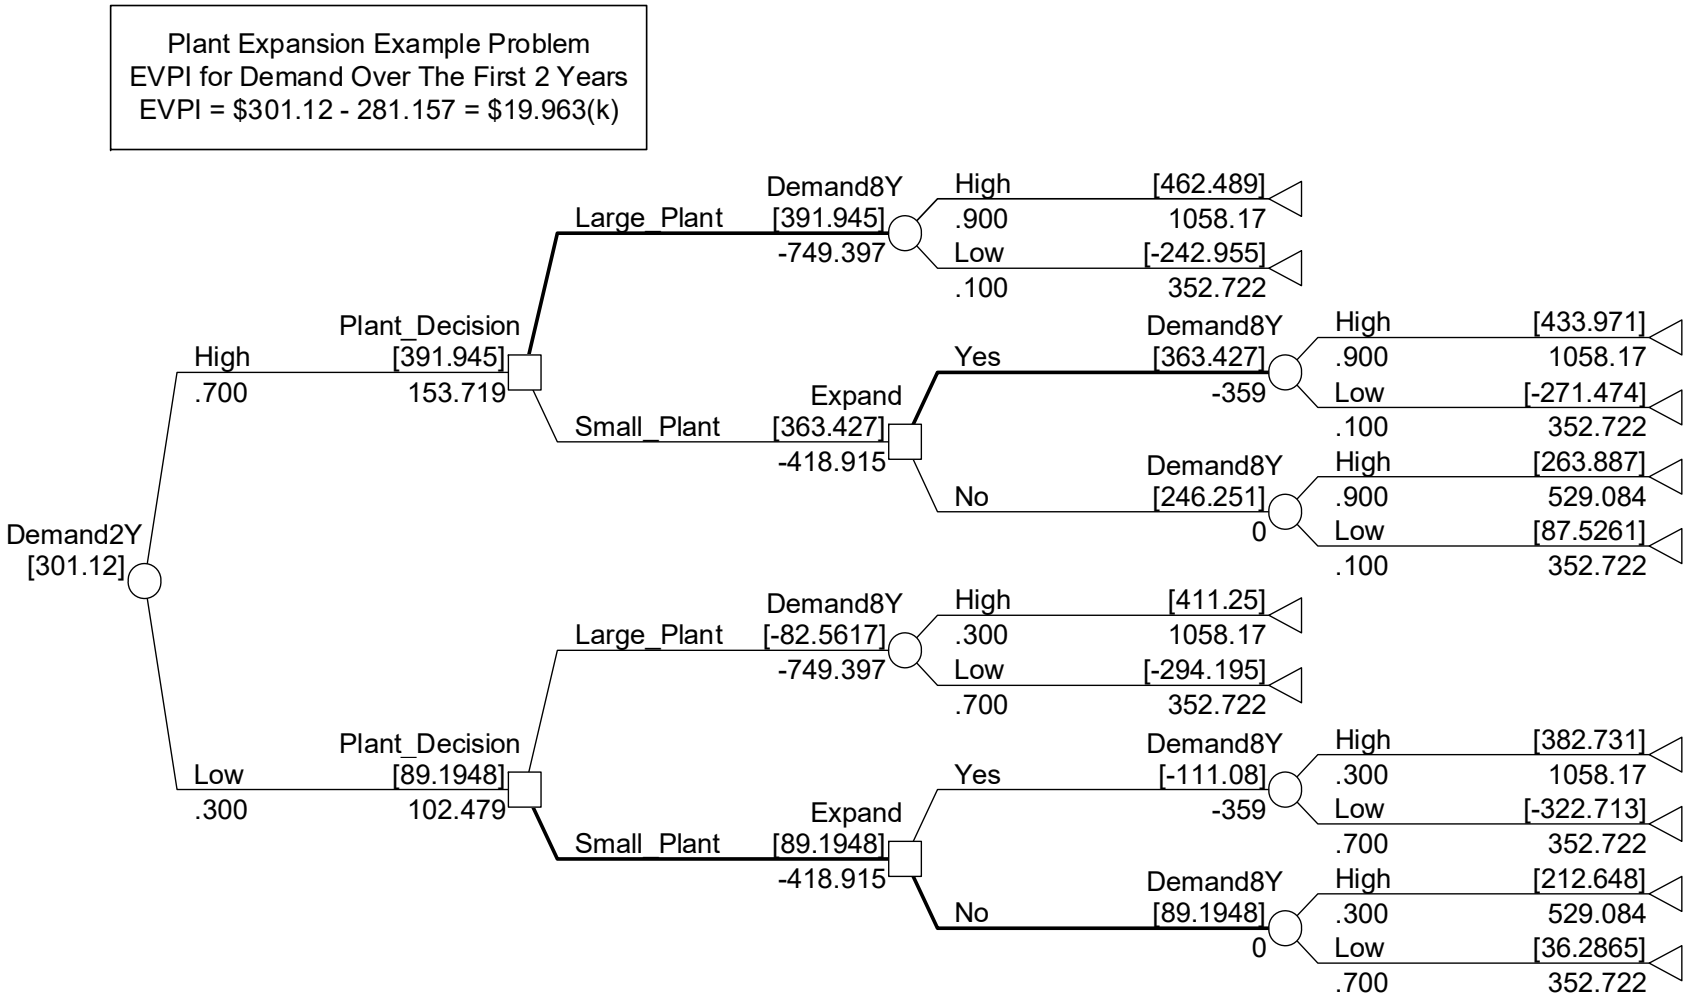


Figure E.24. EVPI about Demand Over First Two Years (DPL)

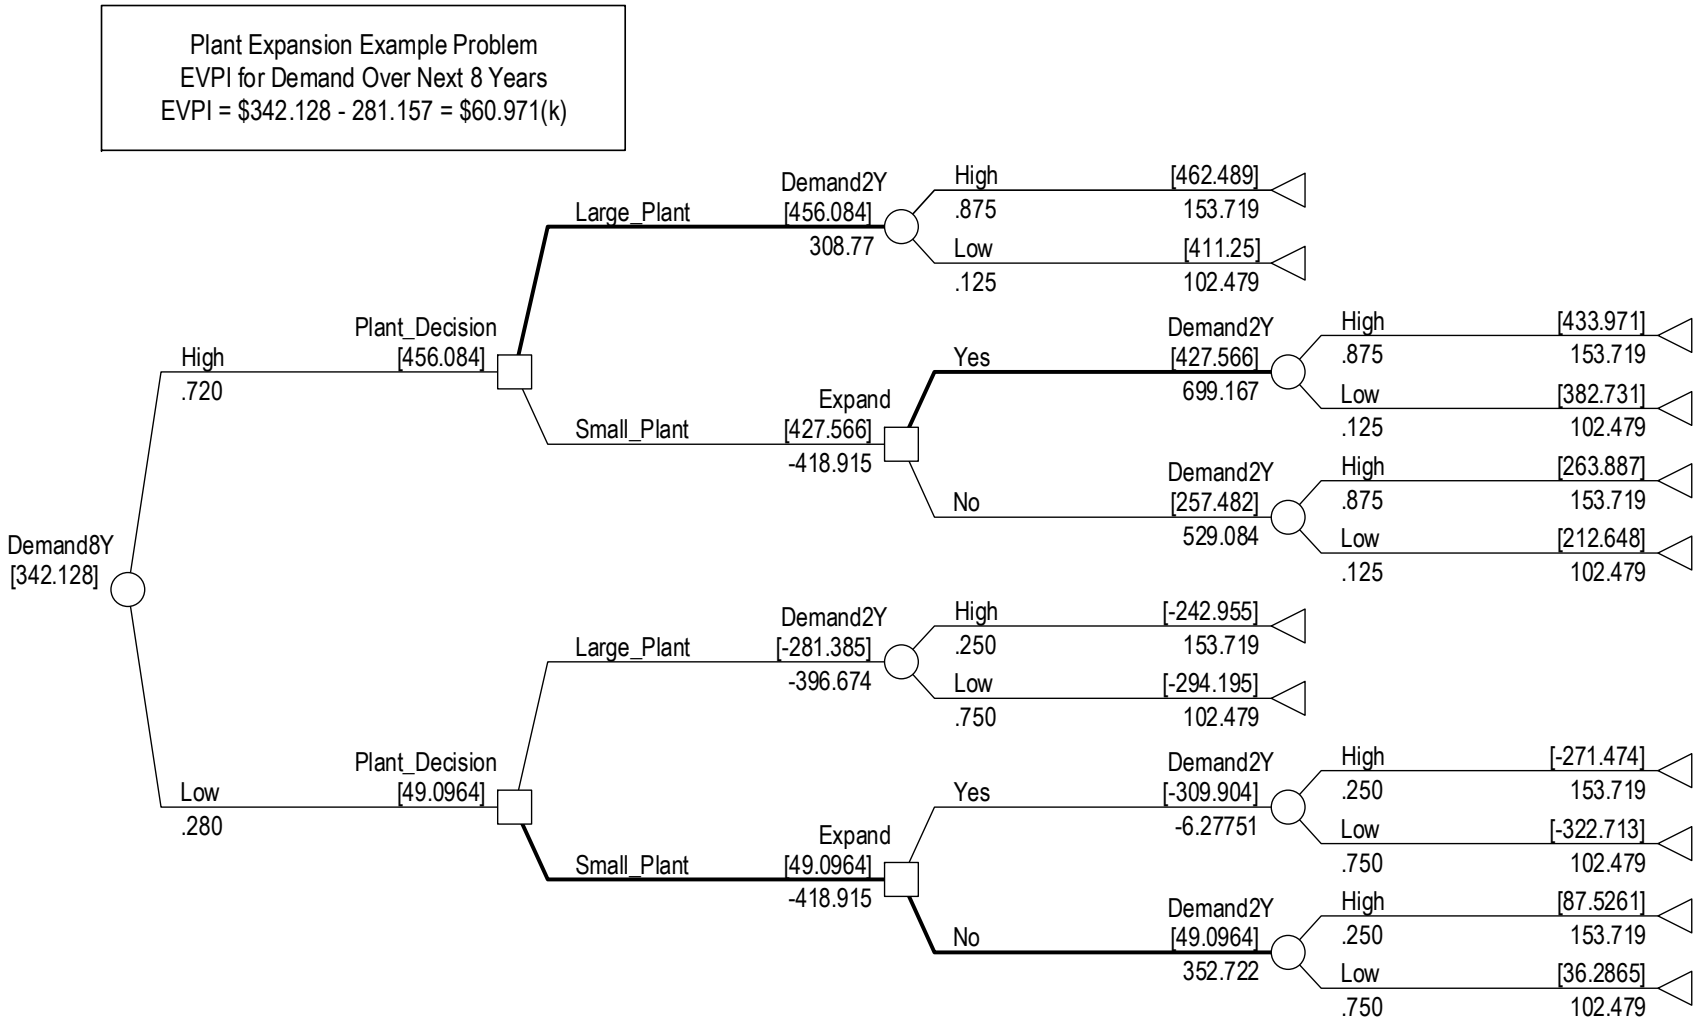
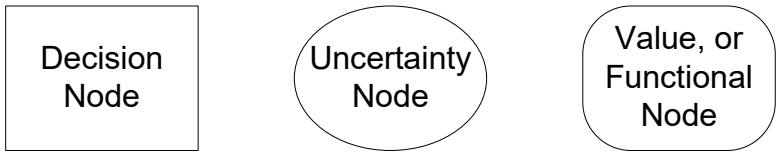


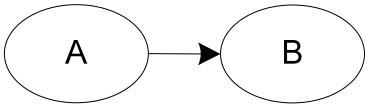
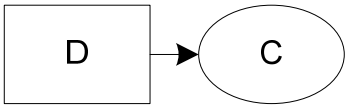
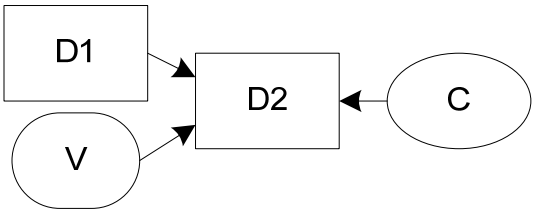
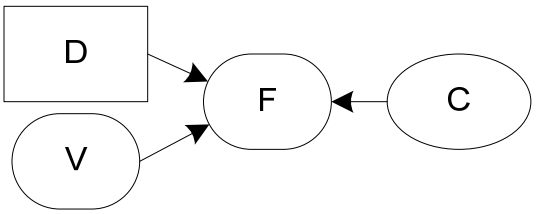
Figure E.25. EVPI about Demand Over Next Eight Years (DPL)

Influence Diagrams



Influence Diagram Nodes

Types of Arrows in an Influence Diagram

 Relevance Arrows	A *relevance arrow* connects two uncertainties *A* and *B* to indicate that P[*B*|*A*] *may* be different from P[*B*]. The absence of an arrow implies lack of relevance: P[*B*|*A*] = P[*B*] (independence)
 Influence Arrows	An *influence arrow* enters an uncertainty *C* from a decision node *D* to indicate that P[*C*|*D*] *depends* on which alternative *D* is chosen. The absence of an arrow implies lack of influence (independence): P[*C*|*D*] = P[*C*].
 Informational Arrows	*Information arrows* enter a decision node *D*2 from another decision node *D*1, or a value node *V*, or an uncertainty node *C*, to indicate that the state of the preceding node *is known* when making the decision *D*2. I.e., arrow *C*-*D*2 means that the outcome of the uncertainty *C* is known before making decision *D*2. It implies *perfect information* about *C*. Similarly, the arrow *D*1-*D*2 means that the choice for *D*1 is known when making decision *D*2.
 Functional Arrows	*Functional arrows* enter a value or functional node (deterministic node) *F* from another value node *V*, or a decision node *D*, or an uncertainty node *C*, to indicate that *F* is a function of its preceding nodes (its inputs).

Expected Value of Sampled (Imperfect) Information (EVSI)

Most research (*information gathering*) available in practice is imperfect because its outcomes do not predict perfectly the outcomes of *state variables*. Such research may reduce but will not eliminate uncertainty about the outcomes of chance nodes. Research may include such actions as model analyses, opinion surveys, purchases of expert advice, site visits, discussions with people, investigations of previous construction, digging test pits, studying literature, or performing tests or preliminary construction tasks. The purpose of research is the reduction of uncertainty of events which follow a decision. The key to the evaluation of research is its effect upon the estimated probabilities of events. The alternative to perform research is shown as just another decision branch on a decision tree that has an associated cost and leads to uncertain research results. The decision on whether to perform research is made by comparing the actual cost of the research to the increase in expected value due to the information provided by the research.

The *expected value of sampled information* (EVSI) is the difference between the expected value of performing the research (assuming it has no cost) and the expected value of the best available option without the research. Because of the assumption that research has no cost, it follows that EVSI >= $0. Obviously, the research should not be undertaken if its cost (CR) exceeds the EVSI. A research alternative is worth doing if it costs less than the increase in expected value it produces, EVSI>=CR. Different research programs will have different EVSI and different costs (CR). The best choice among competing research alternatives is the one with the max(EVSI-CR).

Notice that the cost of the research, CR, is determined by the *market* (i.e., by supply and demand) and has nothing to do with the particular decision problem at hand. In contrast, the EVSI is tied directly to the current decision situation and the present state of information of the decision maker and has nothing to do with how testing is priced in the open market (unless, of course, the provider of the research finds out the EVSI and prices the research accordingly).

Example. A contractor needs a source of aggregate for a project that is somewhat remote and would like to determine the sufficiency of a pit to satisfy the needs of the project. There is a cost of $400,000 to develop the pit and build the facility to process the aggregate, regardless of the quantity of rock in the pit. Test pits can be excavated to estimate better the quantity in the pit, but at a cost. If the pit is developed and does not contain sufficient aggregate, then additional aggregate must be hauled-in to complete the work. The decision alternatives are (a) to develop pit and haul as needed, (b) buy from other suppliers, or (c) perform research before making the decision to develop or buy. Figs. E.26a and E.26b show the complete decision trees. The cost of the test (CPT or CT) is shown as a part of the value of events following the test. Expected value of "development" is -$916.7k which is better than -$1150k for "buy from other suppliers".

Fig E.26a shows that perfect information leads to EMV = -$866.7k, and to EVPI = -$866.7k + $916.7 = $50k. Thus, no imperfect test can be worth more than $50k. Fig E.26b shows that the imperfect test leads to EMV = -$895k. Thus, this test could be performed if its cost CT in the market is less than the EVSI = -$895k + 9l6.7k = $2l.7k. A more (less) accurate test would have a higher (lower) EVSI. In Figs. E.26a and E.26b, we see that if a test was performed, the best aggregate alternative would depend upon its *indication* of *sufficient* or *insufficient* aggregate.

Probabilities on Nature's Trees

A *nature's tree* is an arrangement of chance nodes, their events, and the conditional probabilities for their order in the tree. If chance node C follows chance node A then the probabilities shown on the tree for c_i are conditional upon the occurrence of a_j. Thus, the probabilities shown for C are the conditional $P[c_i \mid a_j]$. The relative positions of chance nodes on a nature's tree can be changed creating another nature's tree. The different order of events on the new tree represents a different set of conditional probabilities. In practice, probabilities can be *estimated* directly for one order of chance nodes on a nature's tree, and the probabilities for a different order of chance nodes on a tree can be *calculated*. The method for calculating one set of conditional probabilities from another is based on *Bayes Theorem* or *Bayes Rule* and can be outlined as follows.

1. Calculate the joint probabilities (probabilities for the intersections) of all events by

$$P[a_i b_j] = P[a_i \mid b_j]P[b_j]$$

2. Calculate the probabilities required for the tree by the *Total Probability Theorem*,

$$P[a_i] = \sum_j P[a_i b_j] = \sum_j P[a_i \mid b_j]P[b_j]$$

$$P[b_j \mid a_i] = \frac{P[a_i b_j]}{P[a_i]}$$

When this last equation is expressed as follows it is known as *Bayes Theorem* or *Bayes Rule*:

$$P[b_j \mid a_i] = \frac{P[a_i \mid b_j]P[b_j]}{\sum_j P[a_i \mid b_j]P[b_j]}$$

Bayes Theorem or *Bayes Rule* provides the means of reversing the order of conditional events in probability statements, i.e. from $P[a_i \mid b_j]$ we derive $P[b_j \mid a_i]$.

It is obvious from these equations that probabilities of any event, including its conditional probabilities, can be calculated from its probabilities of intersection or joint probabilities.

Often, the conditional probabilities that are available (or that are easiest to estimate) are not those required by the order of chance nodes in the decision tree. This is not a problem, however. From the conditional probabilities of events, we can calculate the joint probabilities, and from the joint probabilities we can calculate the probabilities for any nature's tree that has a different order of chance nodes.

This is illustrated by the examples that follow.

© 2024 *Photios G. Ioannou*

Example. The decision tree for the development of the aggregate pit (Fig. E.26a and E.26b had probabilities placed on it for evaluation. The probabilities estimated were: the probability that the aggregate is sufficient is 2/3. If the aggregate is sufficient a test pit will so indicate with probability of 0.9, and if it is not sufficient, the test will so indicate with probability of 0.7. Nature's tree for these probabilities is shown in Fig. E.27a.

Let us define the following events:

S = sufficient aggregate in pit

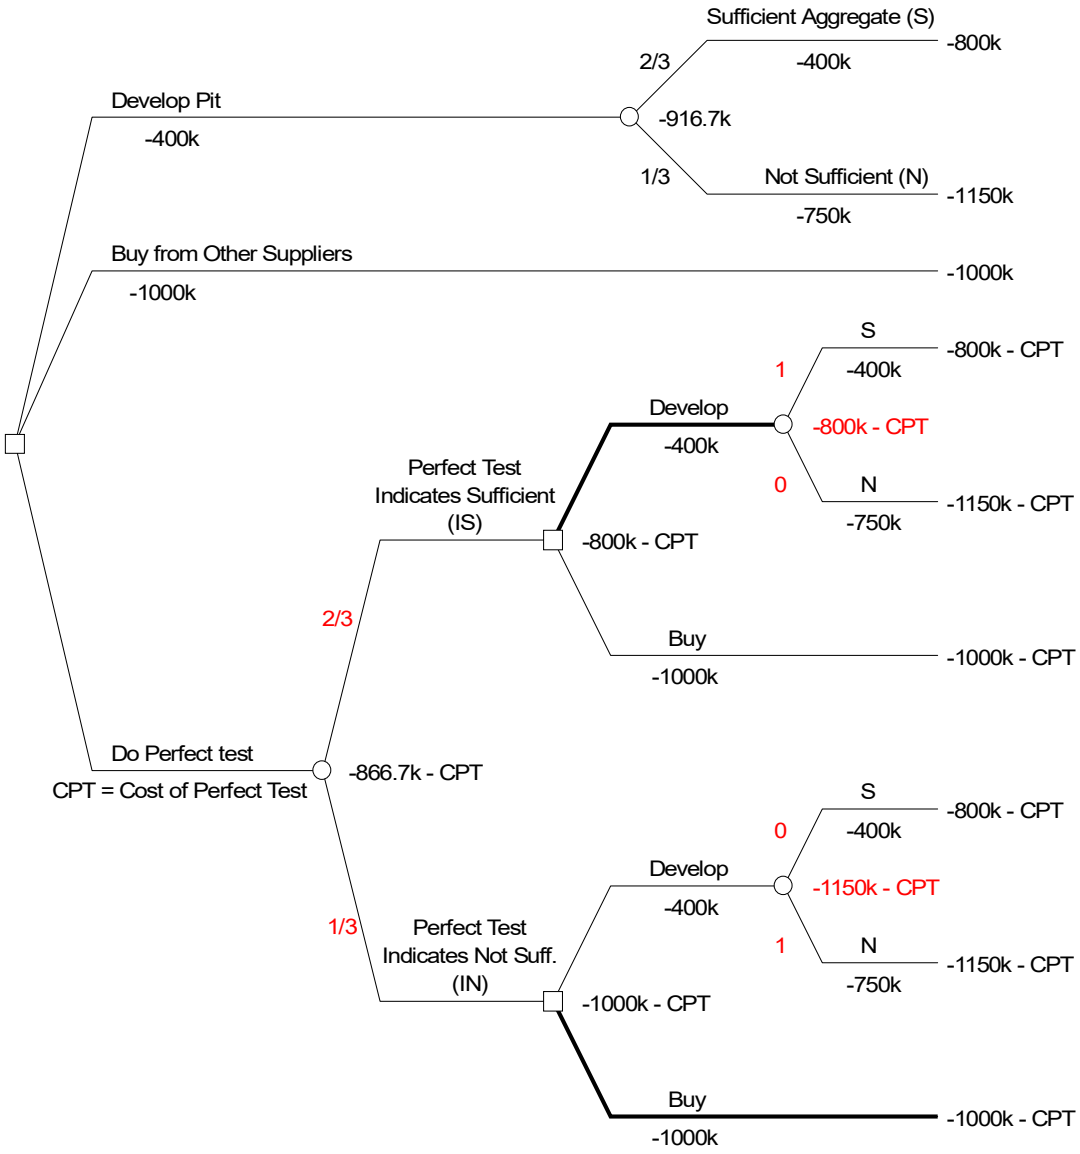
N = not sufficient aggregate in pit

IS = test indicates that there is sufficient aggregate

IN = test indicates that there is not sufficient aggregate

The estimated probabilities have been converted to the form needed for the decision tree, as shown in Fig. E.27b.

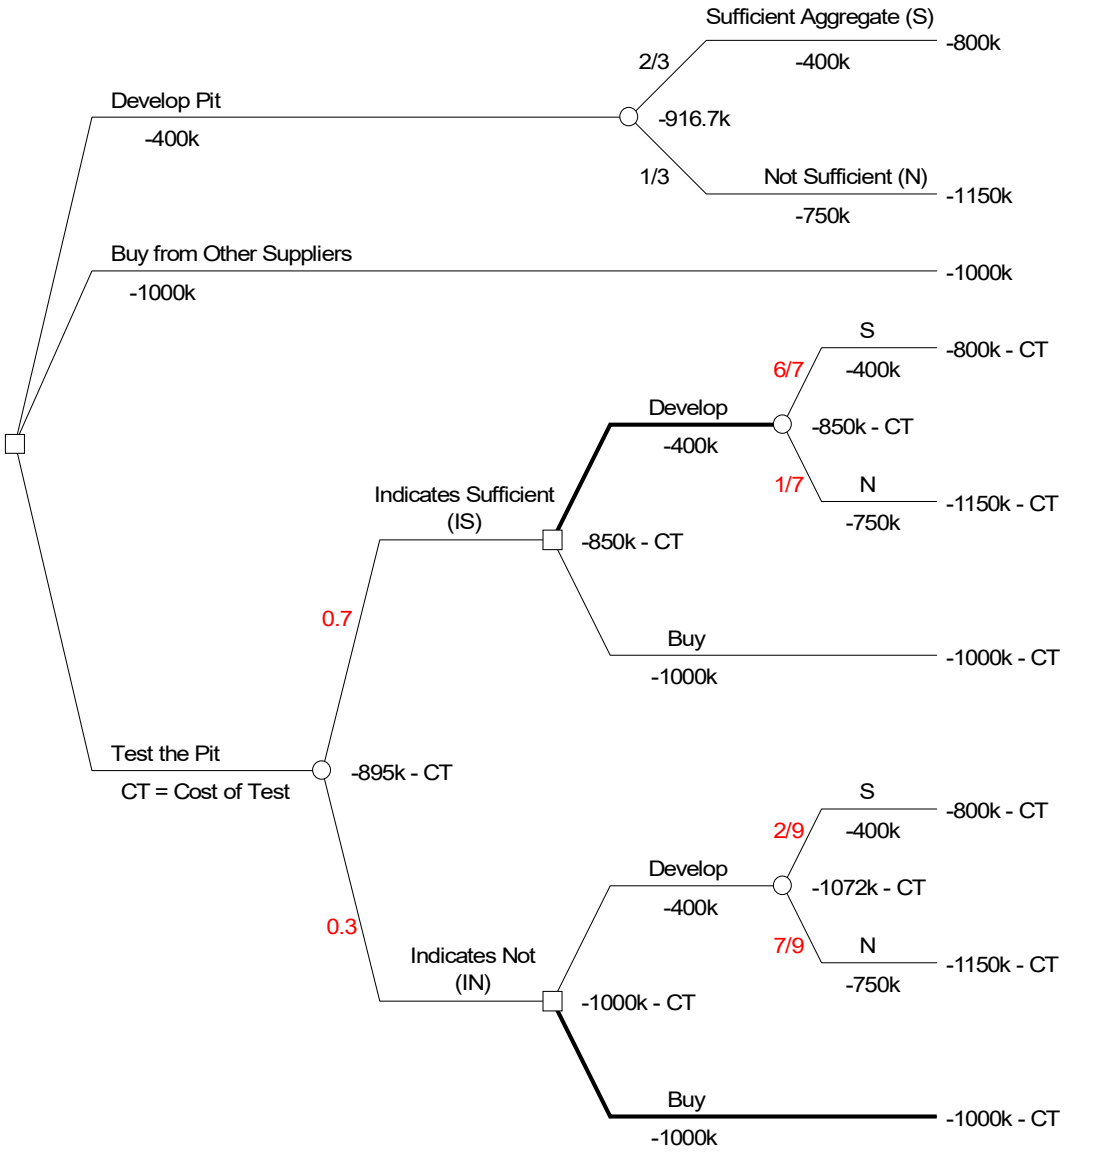
The general conversion of nature's trees for events is shown in Fig. E.30. When larger sets of events are involved, an array structure for conversion, as shown in Fig. E.31 can be useful.



Expected Value of Perfect Information (EVPI)

EVPI = (Expected Value given Perfect Information) - (Max Expected Value without Perfect Information).
For the case above: EVPI = -866.7k + 916.7k = $50k
The EVPI = $50k is an upper bound for the EVSI of any imperfect test about the same uncertainty.
A perfect test should be done only if its cost, CPT, is less than the EVPI. The cost of a perfect test, CPT, is determined by supply and demand in the open market (if available) and not by the decision at hand. In contrast, the EVPI is a function of the test's use to help make a particular decision and thus depends on the test being perfect, the state of information of the decision maker, and the decision's outcomes.
Important: The decision tree for calculating the EPVI of a perfect test has the same structure as the decision tree for calculating the EVSI of any other imperfect test. The difference lies in (a) that the probabilities for observing the perfect test outcomes equal "the prior probabilities of state" (e.g., 2/3 and 1/3) and (b) that the posterior probabilities of state after a perfect test can only be 1 and 0.

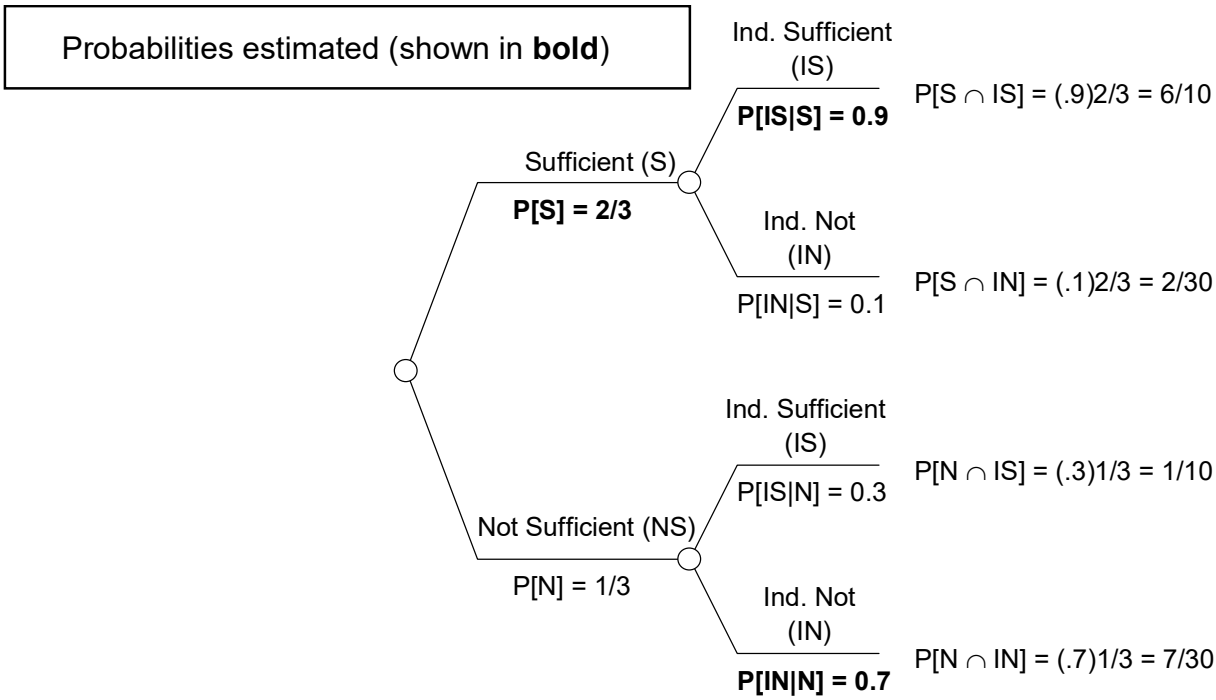
Figure E.26a. Decision Tree for Development of Aggregate Pit and EVPI



Expected Value of Sampled Information (EVSI)

EVSI = (Expected Value given Sampled Information) - (Max Expected Value without Sampled Information)
For the case above: EVSI = -895k + 916.7k = $21.7k
The test should be done only if its cost, CT, is less than the EVSI. Notice that the cost of the test, CT, is determined by supply and demand in the open market and not by the decision problem at hand. In contrast, the EVSI is a function of the test's use to help make this particular decision and thus depends on the test's reliability, the state of information of the decision maker, and the decision's outcomes.
The EVPI for this problem is $50k. This is an upper limit to the EVSI of this and any other any imperfect test. Notice that for this test, the EVSI = $21.7k is about 40% of the EVPI = $50k. It is not uncommon for the EVSI for an imperfect test to be significantly less than the EVPI for the same uncertainty.

Figure E.26b. Decision Tree for Development of Aggregate Pit and EVSI



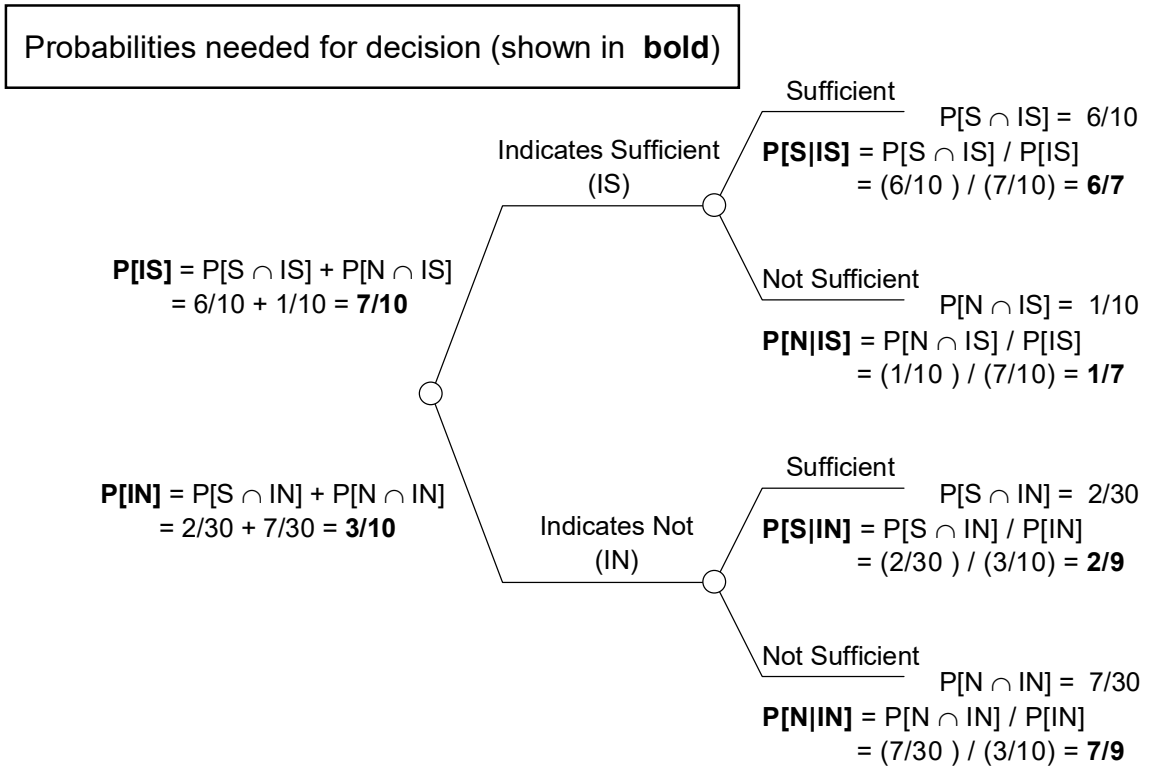
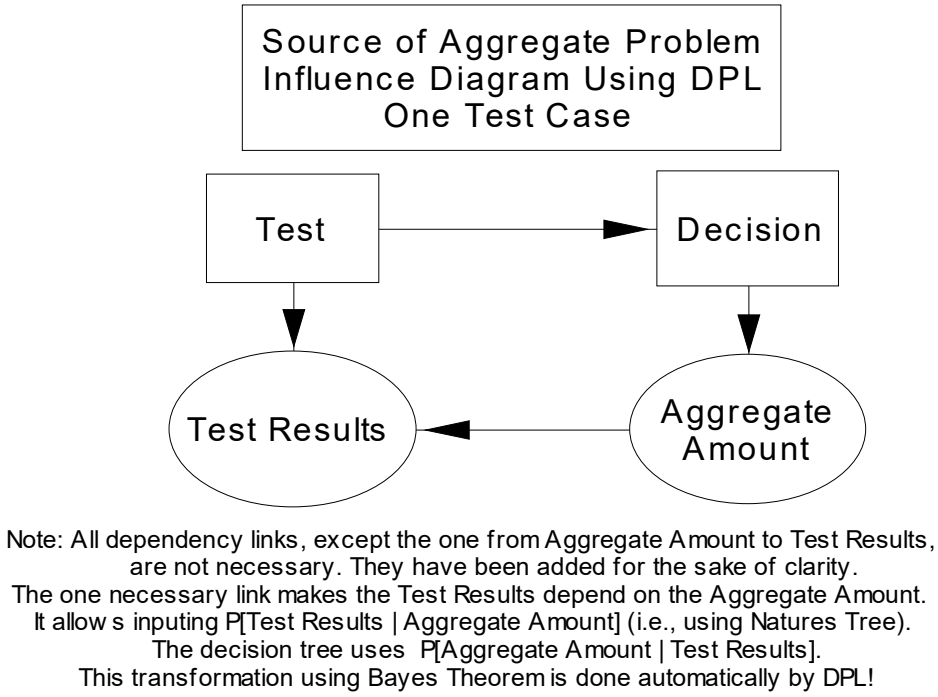


Figure E.27a Reliability of Test and Figure E.27b Probabilities for Decision



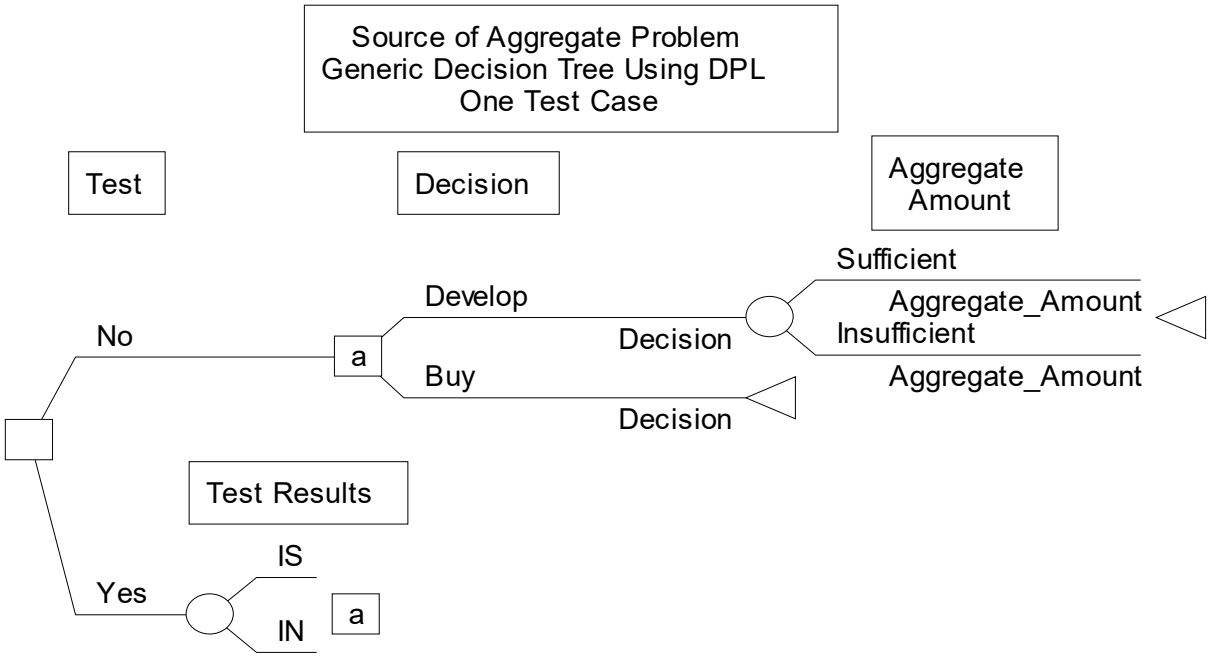


Figure E.28. Generic DPL Decision Tree for Development of Aggregate Pit

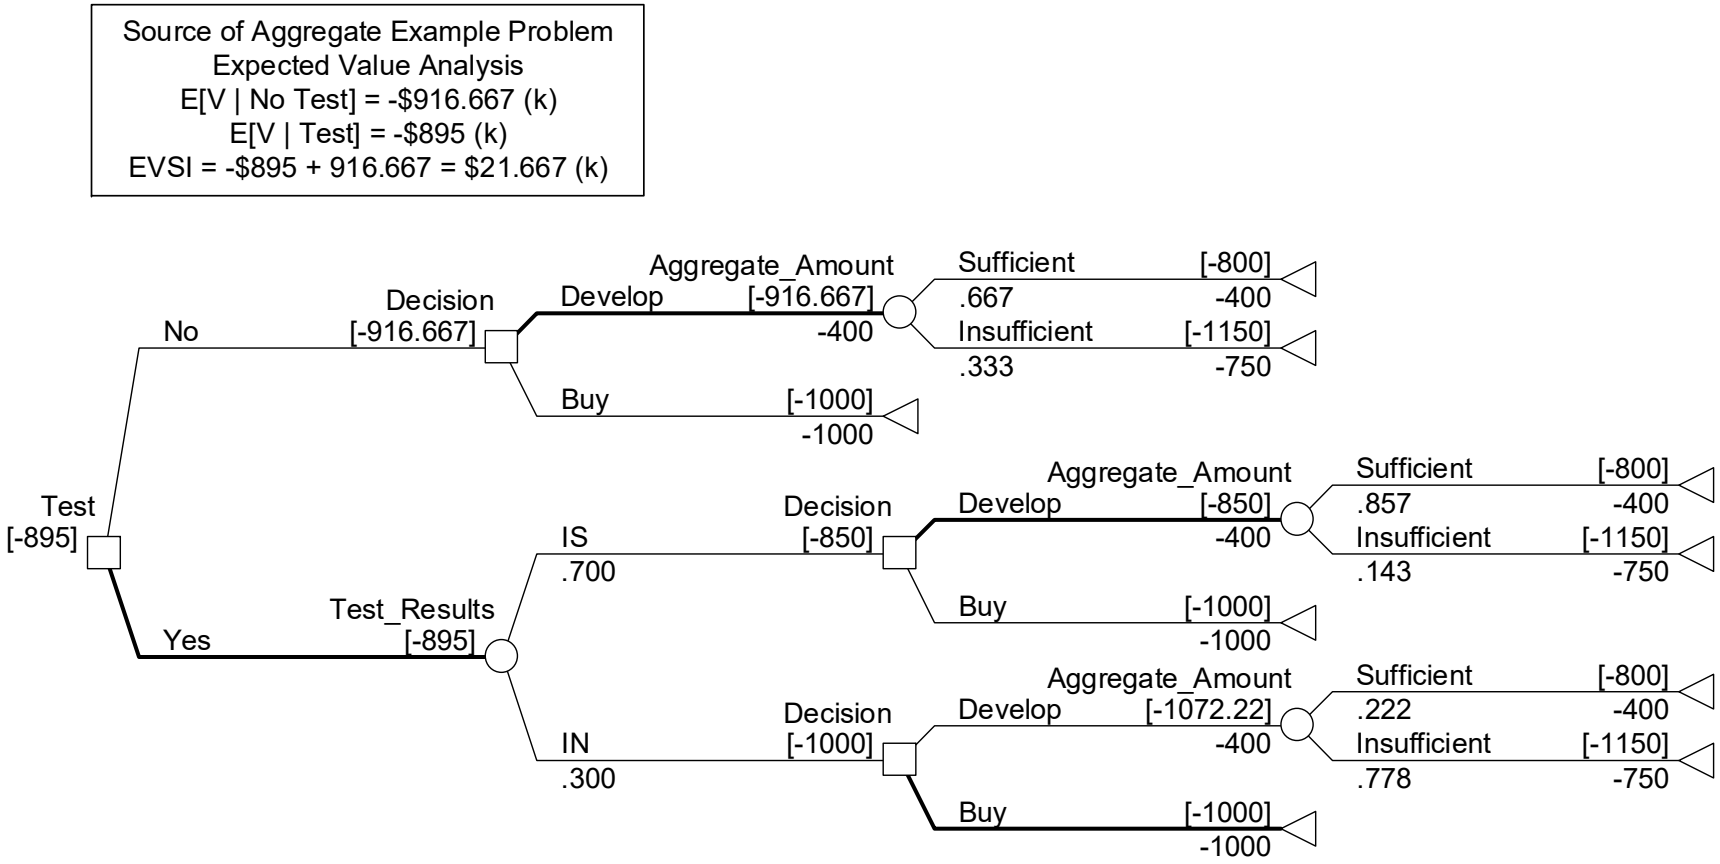


Figure E.29.DPL Model for Development of Aggregate Pit

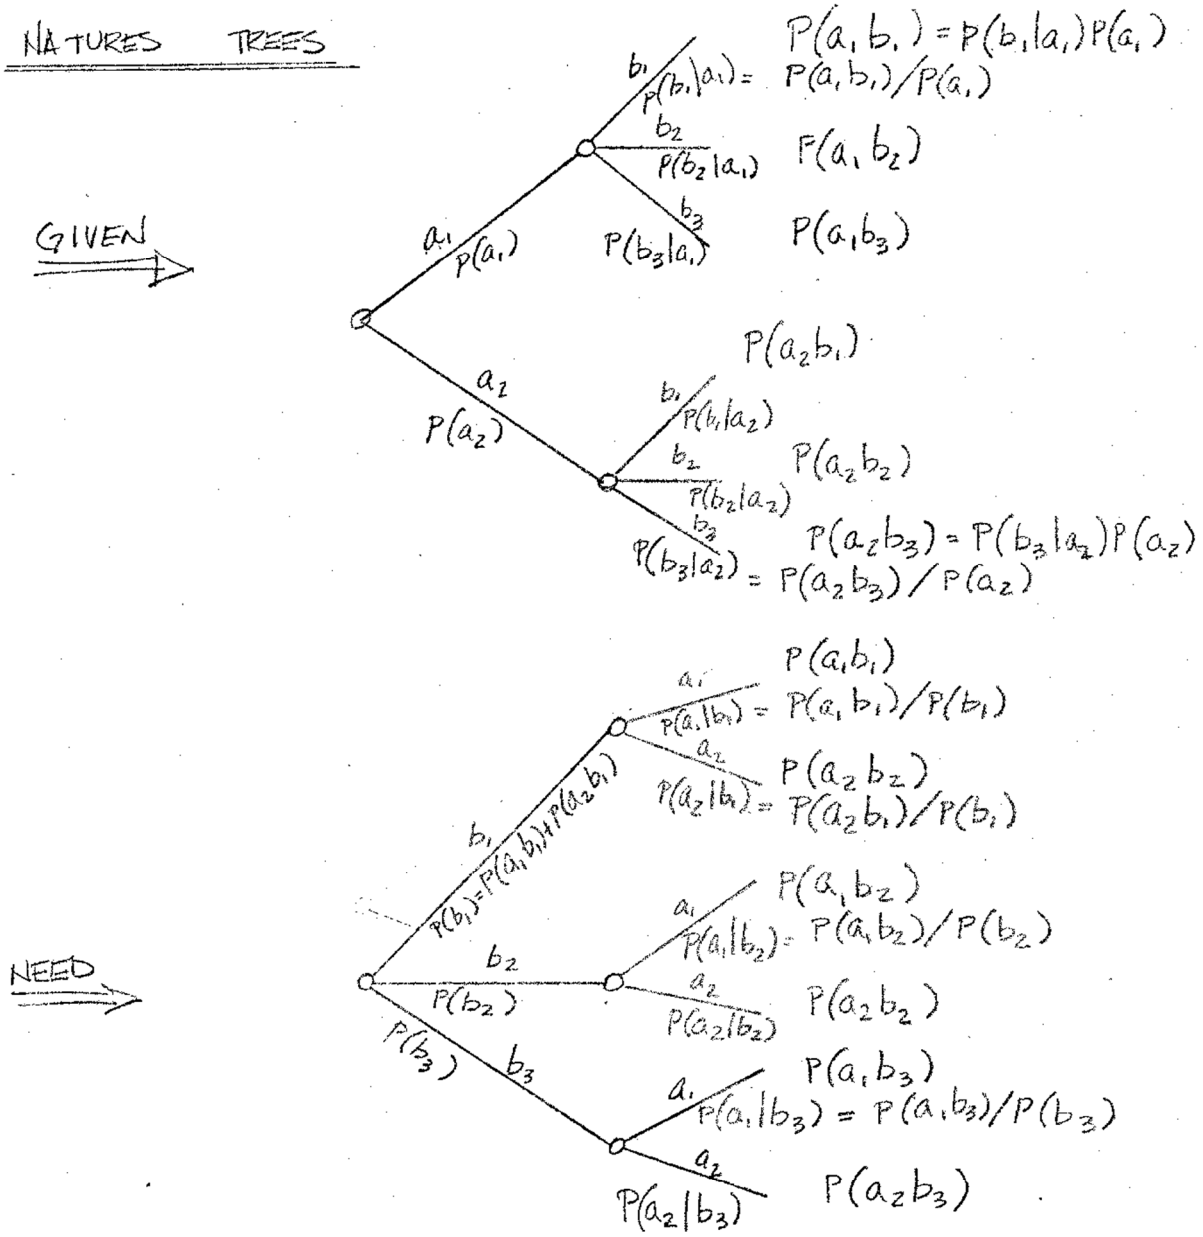
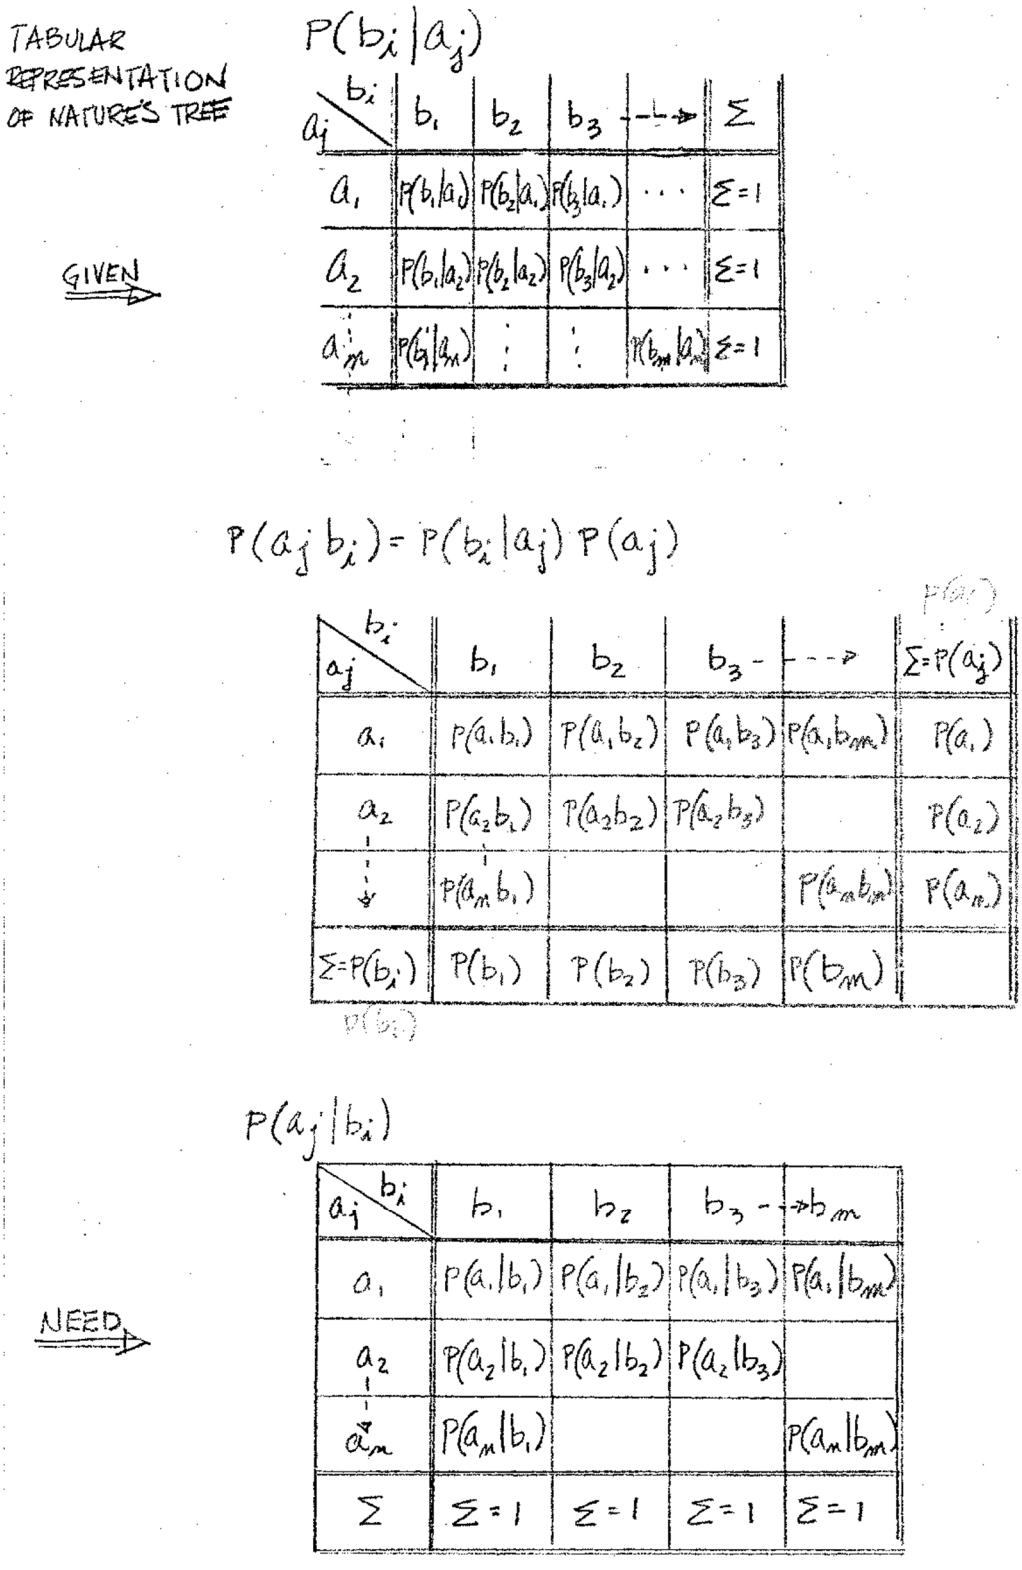


Figure E.30. Reversing Conditional Probabilities Using Event Trees



$P(b_i|a_j)$

a_j \ b_i	b_1	b_2	b_3	⋯	Σ
a_1	$P(b_1\|a_1)$	$P(b_2\|a_1)$	$P(b_3\|a_1)$	⋯	$\Sigma=1$
a_2	$P(b_1\|a_2)$	$P(b_2\|a_2)$	$P(b_3\|a_2)$	⋯	$\Sigma=1$
a_n	$P(b_1\|a_n)$	⋮	⋮	$P(b_m\|a_n)$	$\Sigma=1$

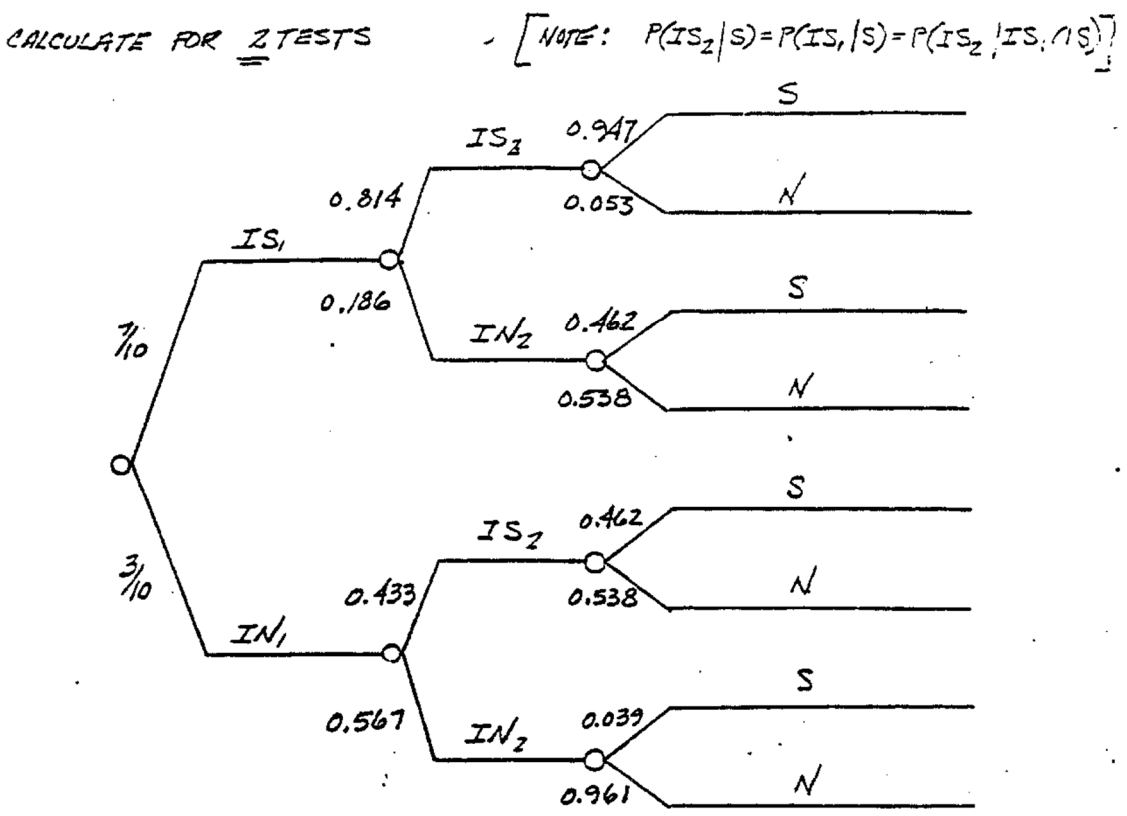
$$P(a_j b_i) = P(b_i|a_j)\,P(a_j)$$

a_j \ b_i	b_1	b_2	b_3	⋯	$\Sigma=P(a_j)$
a_1	$P(a_1b_1)$	$P(a_1b_2)$	$P(a_1b_3)$	$P(a_1b_m)$	$P(a_1)$
a_2	$P(a_2b_1)$	$P(a_2b_2)$	$P(a_2b_3)$		$P(a_2)$
⋮	$P(a_nb_1)$			$P(a_nb_m)$	$P(a_n)$
$\Sigma=P(b_i)$	$P(b_1)$	$P(b_2)$	$P(b_3)$	$P(b_m)$	

$P(a_j|b_i)$

a_j \ b_i	b_1	b_2	b_3	⋯ b_m
a_1	$P(a_1\|b_1)$	$P(a_1\|b_2)$	$P(a_1\|b_3)$	$P(a_1\|b_m)$
a_2	$P(a_2\|b_1)$	$P(a_2\|b_2)$	$P(a_2\|b_3)$	
a_n	$P(a_n\|b_1)$			$P(a_n\|b_m)$
Σ	$\Sigma=1$	$\Sigma=1$	$\Sigma=1$	$\Sigma=1$

Figure E.31. Reversing Conditional Probabilities Using Tables



$$P(IS_2|IS_1) = \frac{P(IS_1 \cap IS_2)}{P(IS_1)} = \frac{P(IS_1 \cap IS_2 \cap S) + P(IS_1 \cap IS_2 \cap N)}{P(IS_1)}$$

$$= \frac{P(IS_1 \cap IS_2|S)P(S) + P(IS_1 \cap IS_2|N)P(N)}{P(IS_1)}$$

$$= \frac{[P(IS|S)]^2 P(S) + [P(IS|N)]^2 P(N)}{P(IS_1)}$$

$$= \frac{(0.9)^2(2/3) + (0.3)^2(1/3)}{0.7} = \frac{0.570}{0.7} = 0.814$$

$$P(S|IS_1 \cap IS_2) = \frac{P(IS_1 \cap IS_2|S)P(S)}{P(IS_1 \cap IS_2)} = \frac{[P(IS|S)]^2 P(S)}{[P(IS|S)]^2 P(S) + [P(IS|N)]^2 P(N)}$$

$$= \frac{(0.9)^2(2/3)}{0.570} = \frac{0.540}{0.570} = 0.947$$

$$P(S|IS_1 \cap IN_2) = \frac{P(IS_1 \cap IN_2|S)P(S)}{P(IS_1 \cap IN_2)} = \frac{P(IS_1|S)P(IN_2|S)P(S)}{P(IS_1|S)P(IN_2|S)P(S) + P(IS_1|N)P(IN_2|N)P(N)}$$

$$= \frac{(0.9)(0.1)(2/3)}{(0.9)(0.1)(2/3) + (0.3)(0.7)(1/3)} = \frac{0.06}{0.13} = 0.462$$

Figure E.32a. Aggregate Pit Example—Conditional Probabilities of State Given Two Tests

Flow of Joint Probabilities on Event Tree and their Split by Conditional Probabilities

The event tree in Figure E.32b shows the event nodes for the *Source of Aggregate* example.

- The branches for the first node on the left show the states of nature "S" or "N" and their prior probabilities, P[S] and P[N].
- The branches for the nodes to the right show the outcomes of the tests Test 1 and Test 2 that follow. These events are conditional on the given true state of nature "S" or "N".
- The branches for the nodes to the right also show the conditional probabilities:
  - P[IS | S], P[IN | S], P[IS | N], P[IN | N] for Test 1 and for Test 2.
- These conditional probabilities are the same as the "Reliability of the Test" shown in Figure E.27a and are the same for both Test 1 and for Test 2 because these conditional events are independent (because they are conditional on the true state of nature "S" or "N" and thus cannot change the probability of the true state of nature "S" or "N").
- Each branch also shows "the joint probability ∩" for the node that follows which is calculated by multiplying "the joint probability "∩" for the preceding branch", times "the conditional probability of that branch".
- For example,
  - P[S, IS1] = P[S] * P[IS1 | S] = 0.667 * 0.9 = 0.600
  - P[S, IS1, IS2] = P[S, IS1] * P[IS2 | S, IS1] = 0.6 * 0.9 = 0.54
- A probability of 1 starts at the root node on the left side of the tree and flows to the right through the nodes and branches like water through pipes.
- The probability flow through each branch is given by the joint probability "∩".
- This probability flow is split at each node. The portion that flows through each succeeding branch is given by the conditional probability of that branch.
- The split flow gives the joint probabilities for the succeeding branches.

In Figure E.32c, the true states of nature "S" or "N", appear at the end of tree, i.e. after the results of Test 1 and Test 2 are known.

- Thus, the probabilities for the true states of nature "S" or "N" are posterior to the preceding test observations and must be computed using Bayes Theorem.
- To apply Bayes Theorem, we copy the joint probabilities from the end branches of the tree in Figure E.32b to the corresponding end branches of the tree in Figure E.32c.
- The colors "red", "purple", "blue", and "green" indicate the corresponding branches between the first and the second tree.
- We then trace and combine the flow of probabilities in the tree, from right to left:
  - For example, starting on the right of the tree in Figure E.32c and moving left, the red joint probabilities combine and give: 0.540 + 0.030 = 0.570.
  - Similarly, when moving left, the purple joint probabilities at the right of the tree in Figure E.32c combine and give: 0.060 + 0.070 = 0.130.
- Moving further left, the red 0.570 and the purple 0.130 joint probabilities combine and give: 0.570 (red) + 0.130 (purple) = 0.700 (red and purple) = P[IS1].
- All joint probabilities "∩" are calculated similarly by combining probability flows.
- The conditional probabilities for the branches are then calculated by dividing "the joint probability of each branch" by "the joint probability of its preceding branch". E.g.:
  - P[IS2 | IS1] = P[IS1, IS2] / P[IS1] = 0.570 / 0.7 = 0.814.
  - P[S | IS1, IS2] = P[IS1, IS2, S] / P[IS1, IS2] = 0.540 / 0.570 = 0.947.

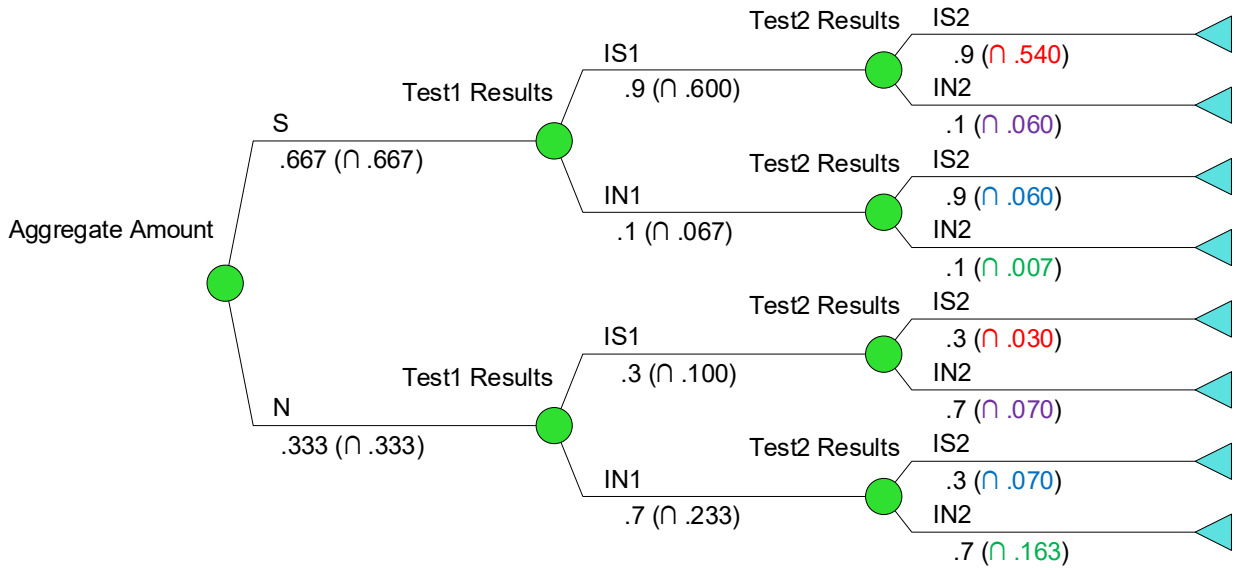


Figure E.32b. Aggregate Pit Example—Prior States of Nature Followed By Two Tests

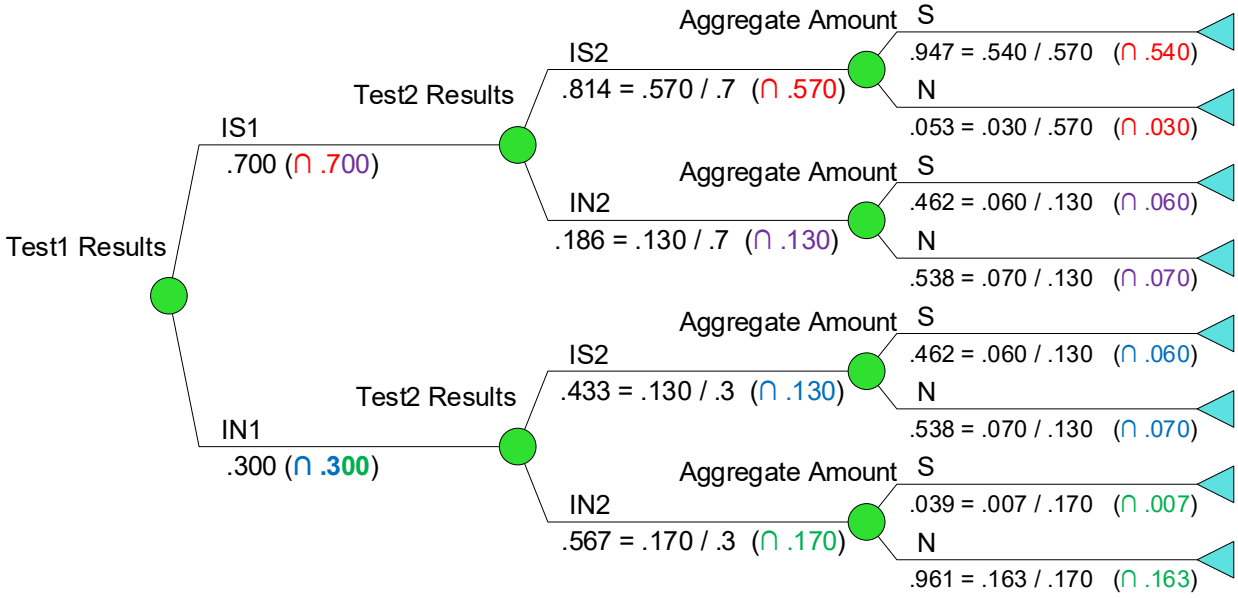


Figure E.32c. Aggregate Pit Example—Conditional Probabilities of State Given Two Tests
(Resulting probabilities are identical to those in the tree shown in Figure E32.a)

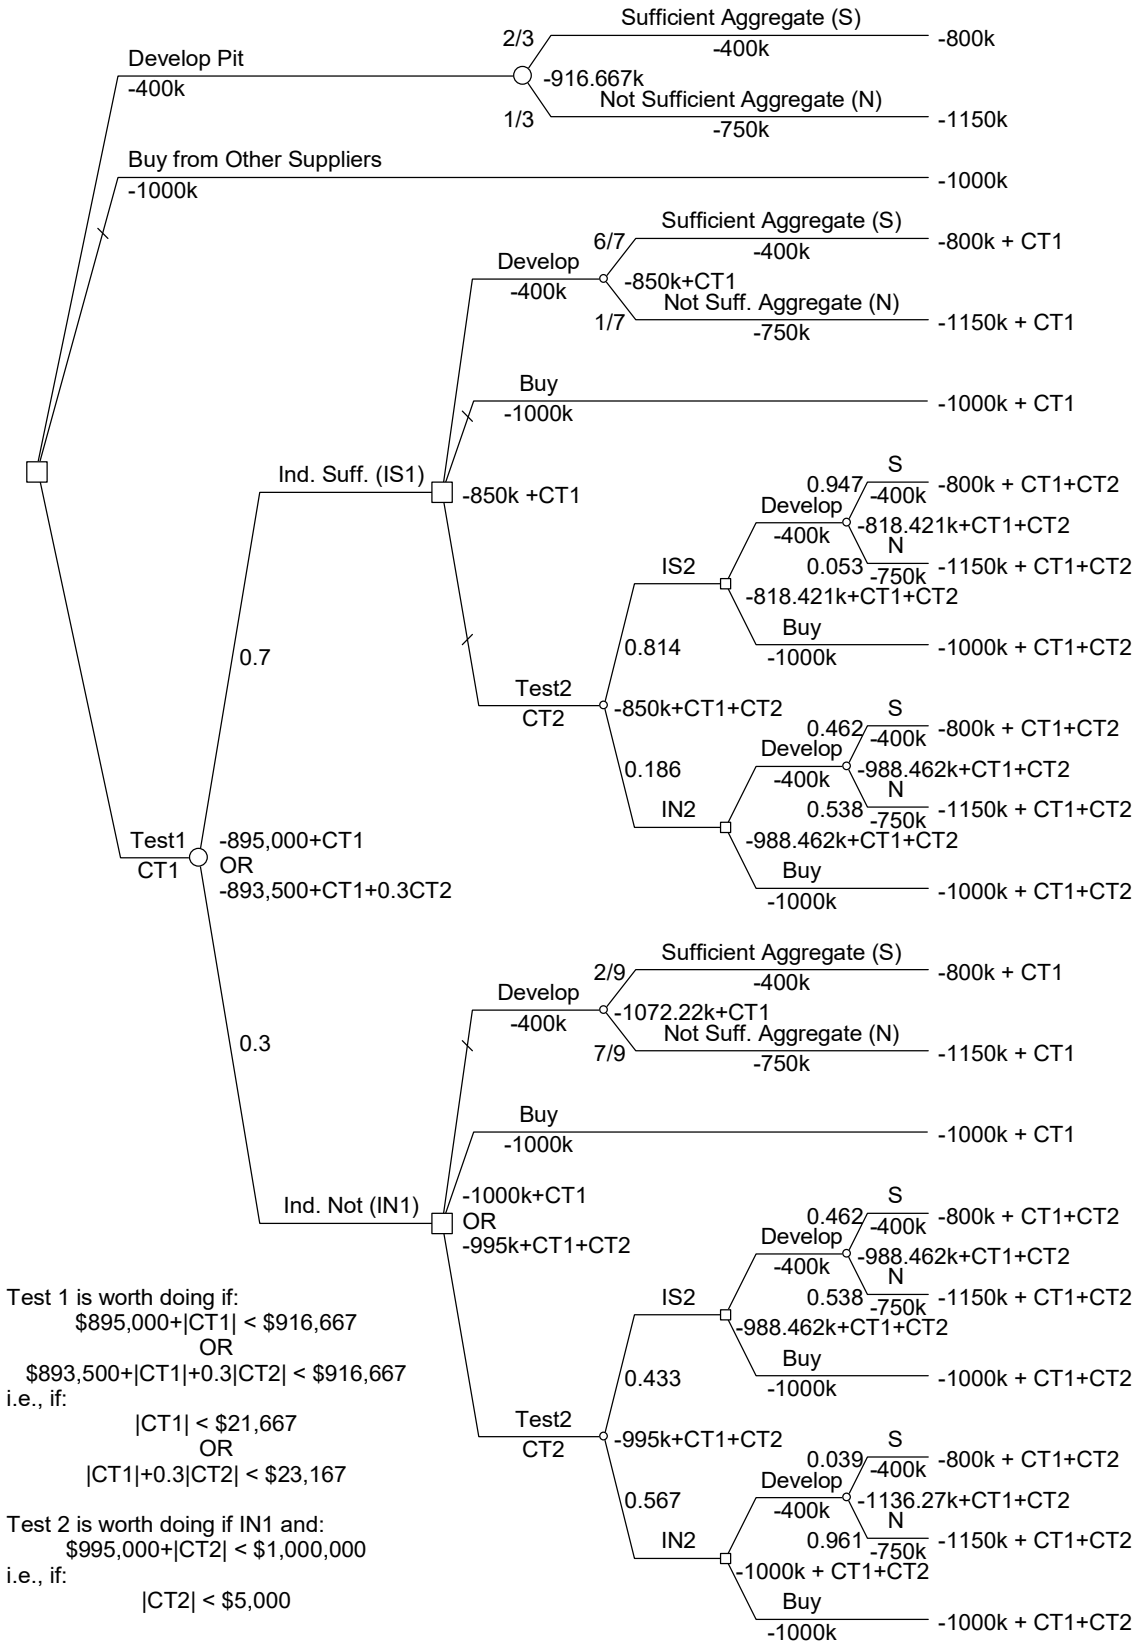
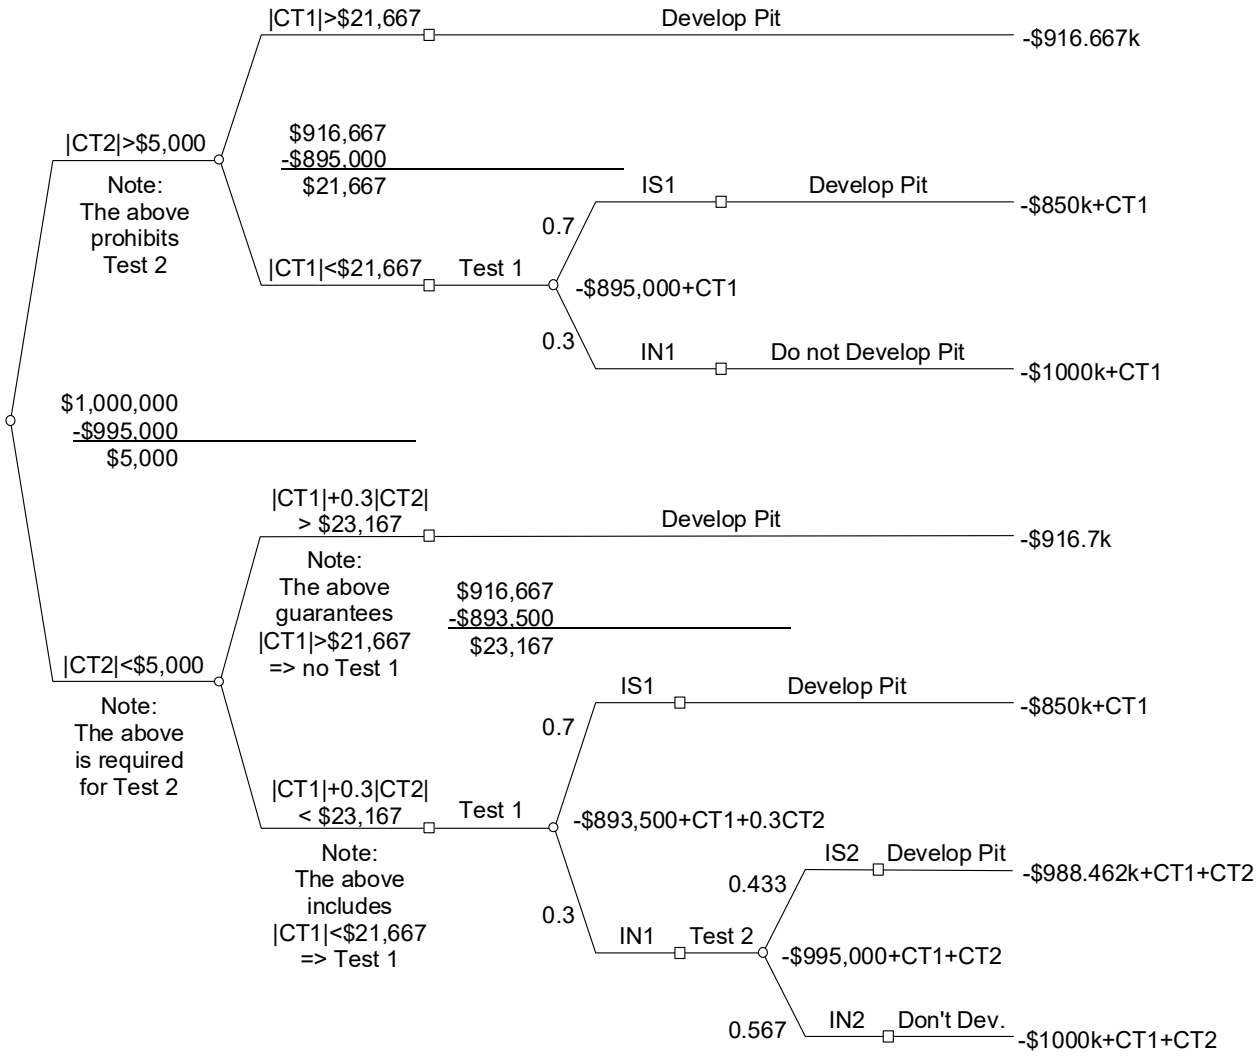


Figure E.33. EVSI for Two Tests for Development of Aggregate Pit



Test2: EVSI2 = $5,000, but only when Test1 shows IN1 (if Test1 shows IS1, EVSI2 = $0). Thus, conduct Test2 only when (Test1 indicates IN1) and (CT2 < $5,000).

Test1: EVSI1 for Test1 depends on the cost of Test2, CT2, as follows:

CT2	EVSI1	Cost Constraint	Do Test2?
$10,000	$21,667	|CT1|<$21,667	No
$5,000	$21,667	|CT1|<$21,667	No
$5,000	$21,667	|CT1|+0.3|CT2|<$23,167	IN1 ? Yes : No
$1,000	$22,867	|CT1|+0.3|CT2|<$23,167	IN1 ? Yes : No
$0	$23,167	|CT1|+0.3|CT2|<$23,167	IN1 ? Yes : No

Note: EVSI2 << EVSI1, i.e., Test1 (being first) is much more valuable than Test2. EVSI1 increases from $21,667 to $23,167, as CT2 decreases from $5,000 to $0. The possibility that Test2 may follow Test1, increases the EVSI1 of Test1.

Figure E.34. Two-Test Policy for Development of Aggregate Pit

The Value of Performing Multiple Tests

Let us consider again the "source of aggregate" problem. Decision trees E26b and E29 can be used to determine the EVSI of performing one test. Decision trees E33, E34, and E36 can be used to determine the EVSI of performing two tests. The EVSI for the two-test case can be investigated with or without the option to stop testing after the first test.

The following expected values were computed using the above decision trees under the assumption that both the first and the second test are free: CT1 = CT2 = 0. Based on these expected values we will consider the EVSI of one test or two tests under the contractual possibilities described below.

EV(No tests) = -$916,667
EV(Test 1 only) = -$895,000
EV(Test 1 followed by Test 2) = -893,500

EV(No Test 2 | 1st Test IS) = -$850,000
EV(Do Test 2 | 1st Test IS) = -$850,000
➔ EVSI(Test 2 | 1st Test IS) = $0

EV(No Test 2 | 1st Test IN) = -$1,000,000
EV(Do Test 2 | 1st Test IN) = -$995,000
➔ EVSI(Test 2 | 1st Test IN) = $5,000

EVSI for Three Contractual Situations for Performing Tests

A. The contractor has to sign a contract today and pay now for all tests that the testing company will perform in the future.

   - A1. Up to how much would the contractor be willing to pay today for the testing company to perform one test?

     EVSI(1 test) = -895,000 + 916,667 = $21,667

   - A2. Up to how much would the contractor be willing to pay today for the testing company to perform two tests?

     EVSI(2 tests) = -893,500 + 916,667 = $23,167

B. The contractor has to sign a testing contract and pay an amount *CT1* today to perform one test with the option to perform a second test at an additional cost *CT2*. The decision whether to pay *CT2* and perform the second test will be made after the outcome of the first test is observed (IS1 or IN1).

   - B1. Suppose that the second test is free, i.e., *CT2*=$0. Up to how much would the contractor be willing to pay for one test with the option of a second test?

     EVSI(Test 1 & Option | CT2=$0k) = -893,500 + 916,667 = $23,167

- B2. Suppose that the second test costs CT2 < 5k. In this case, EVSI2 = -$995+1000 = $5k *only* if the first test IN1 (P[IN1]=0.3)

EVSI(Test 1 & Option | CT2<$5k) = -893,500-0.3CT2+916,667 = $23,167-0.3CT2

- B3. Suppose that the second test costs CT2 > 5k. In this case, the second test is too expensive and should not be performed.

EVSI(Test 1 & Option | CT2>$5k) = -893,500-0.3x5,000+916,667 = $21,667

C. The contractor is offered a contract to pay the testing company an amount *CT1* < EVSI(1 Test) = $21,667 to perform one test. For example, *CT1* could be $20,000.

Before the contract is finalized, the testing company offers the option to perform a second test at an additional cost *CT2* (a specific amount) to be paid after the outcome of the first test is observed (IS1 or IN1). To have this option added to the contract the contractor needs to pay now an additional amount *V*.

- C1. Up to how much *more* than CT1 (e.g., the $20,000) for one test would the contractor be willing to pay today to add the option to perform a second test (at an additional cost CT2)?

- EV(Test 1 only) = -$895,000 – CT1
- EV(Test 1 followed by Test 2 | CT2>$5k) = -895,000 – CT1

EVSI(Option | CT2>$5k) = -895,000–CT1 + 895,000+CT1 = $0

If CT2>5k then *V*=$0

- EV(Test 1 followed by Test 2 | CT2<$5k) = -893,500 – CT1 -0.3CT2

EVSI(Option | CT2<$5k) = -893,500–CT1-0.3CT2 + 895,000+CT1 = $1,500 – 0.3CT2

If CT2<5k then *V*<=$1,500 – 0.3CT2

Notice that the EVSI for the option to perform Test 2 does not depend on the actual cost CT1 of performing Test 1 (as long as the first test will be performed). Thus, the maximum value for *V* does not depend on whether CT1=$20,000 or CT1 =$10,000.

The maximum value of the option, *V* = $1,500, occurs when Test2 is free, CT2=$0.
The value of the option *V* decreases linearly as the cost CT2 increases: *V* = $1,500 – 0.3CT2.
When CT2>=$5,000, then the option to perform Test2 has no value, *V*=$0.

These results can be summarized as follows:

CT2 = $0	*maxV* = $1,500
$0 < CT2 < $5,000	*maxV* = $1,500 – 0.3CT2
CT2 >= $5,000	*maxV* = $0

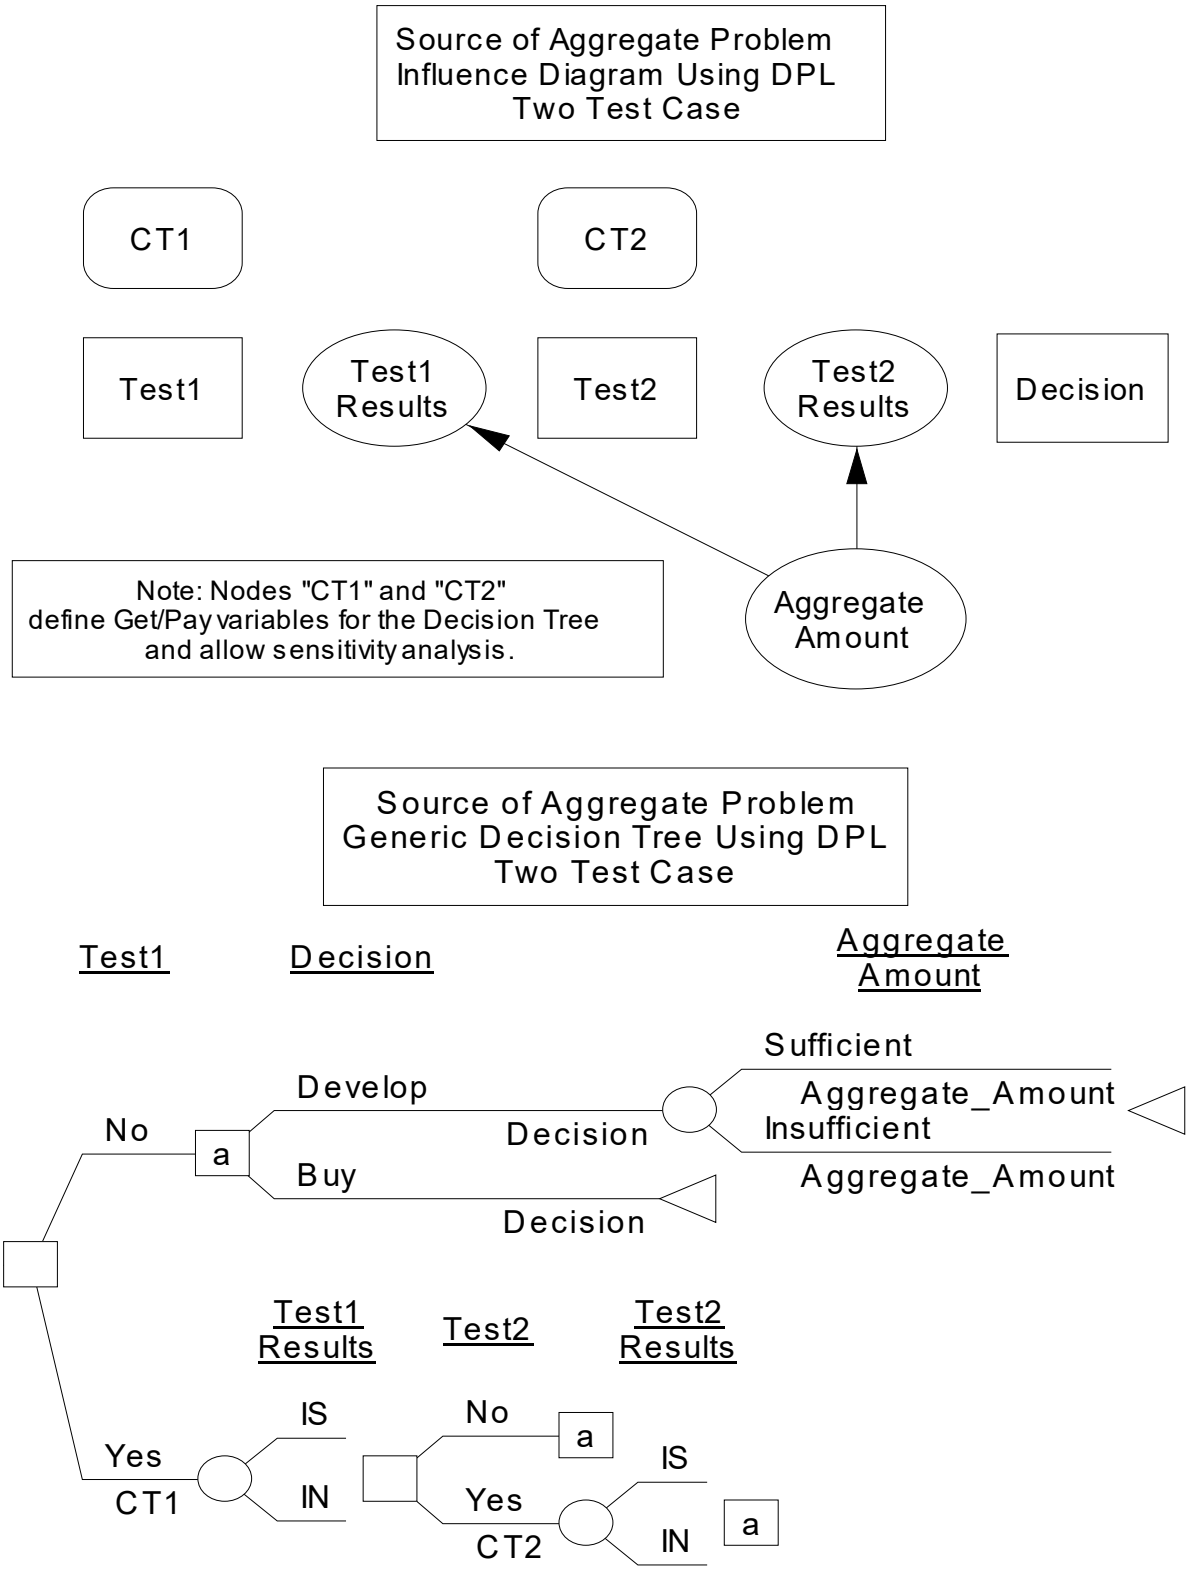


Figure E.35. Generic Two-Test DPL Decision Tree for Development of Aggregate Pit

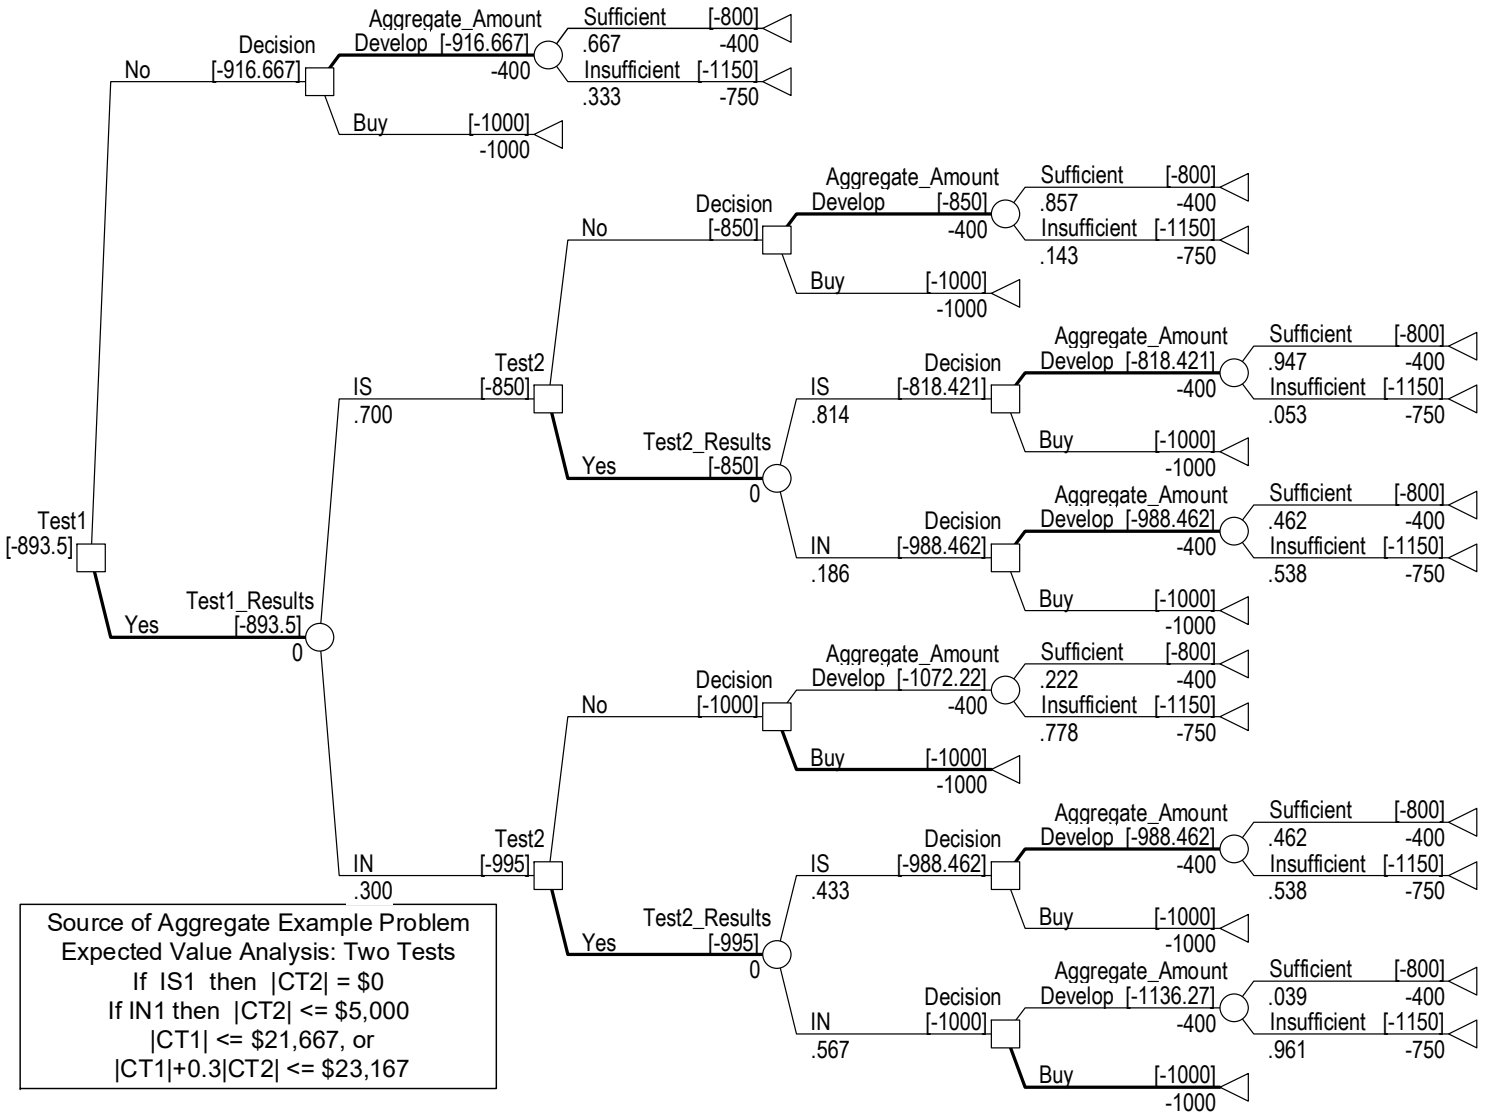


Figure E.36. Two-Test DPL Decision Tree for Development of Aggregate Pit

Relationship between EVSI and EVPI —Automated calculation of EVSI in DPL

The expected value of perfect information (EVPI) is the increase in expected value due to an ideal test that predicts with certainty the outcomes of a particular chance variable.

The expected value of sampled information (EVSI) is the increase in expected value due to a real-world (imperfect) test that can provide information for a particular chance variable.

Clearly, for the same particular chance variable we have: EVPI >= EVSI.

While the above are true, the EVPI and EVSI are also connected in a very important way:

EVSI of an experiment = EVPI on the *observations of the experiment*

This is because

- "*performing a particular experiment and observing its outcome*"

is logically equivalent to

- "*asking a clairvoyant what the outcome of the experiment will be*".

In both cases, we will observe (or, we will be told) the outcome of the experiment *before* we make the decision for which the experiment provides material information.

Automated calculation of EVSI in DPL

DPL can provide automatically the EVPI for any chance variable in a decision model.

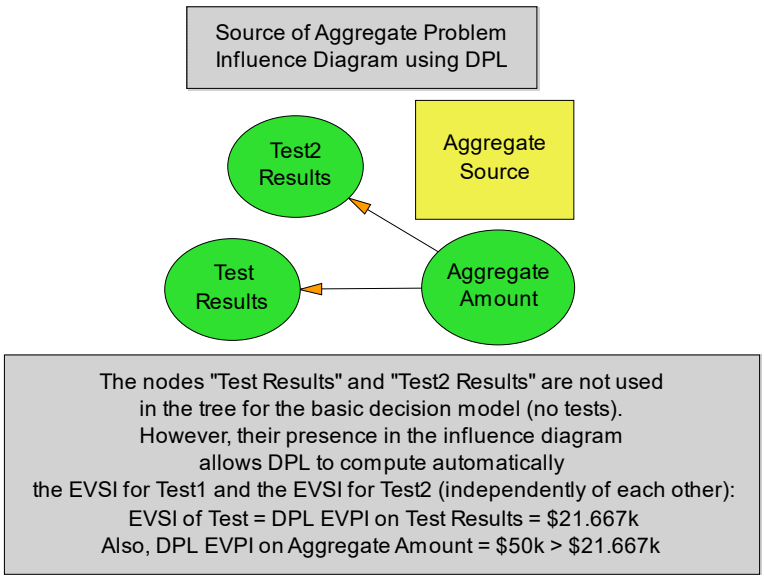
Thus, the above observation makes the automatic computation of EVSI in DPL very easy.

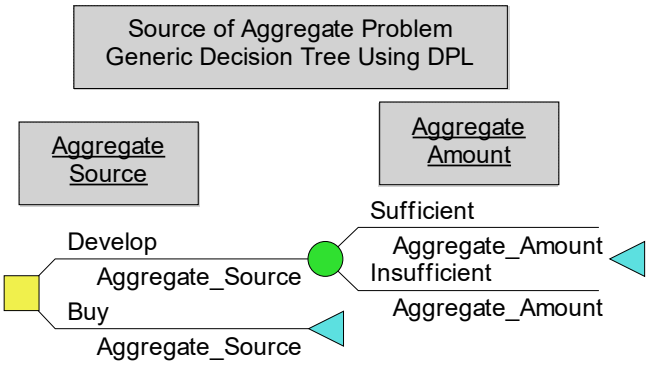
To have DPL compute automatically the EVSI of a test *T* that provides information about chance node *X*, follow the steps below:

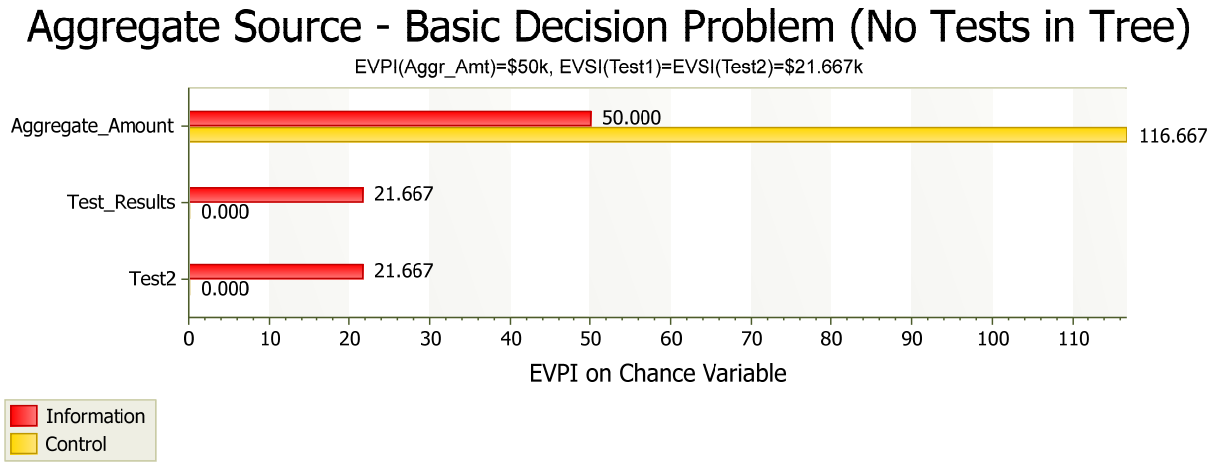
- In the influence diagram for a basic decision model, add a new chance node, *T*, that represents the test results. This new chance node needs to be defined in the influence diagram only and *does not need to be part of the decision tree*.
- In the influence diagram, add an arrow from *X* to *T* to indicate that the probabilities to be entered for *T* are conditional on *X*, i.e., that P[*T*|*X*].
- Enter the conditional probabilities P[*T*|*X*], which are those given directly by the *likelihood probability matrix* or the *reliability* for the test observations. There is no need to use Bayes Theorem and compute P[*X*|*T*].
- In DPL, run a "*decision analysis*" and choose to produce automatically the "*Value of Information & Control*" on all chance variables. The reported "*Value of Information*" (i.e., the EVPI) on the chance variable "*T*" is the EVSI for the test.

This DPL capability provides a very convenient way to obtain the EVSI for any test quickly and easily. However, it is not a complete substitute for the addition of the *test option* to the decision tree and analyzing the test alternative. The advantage of a decision tree with a test option is that it also provides the *optimal policy* (best decision alternative) for each possible test outcome. Nevertheless, even in this case, the EVSI produced automatically by DPL as the "EVPI for *T*", is useful for validating the EVSI calculated manually through the decision tree.

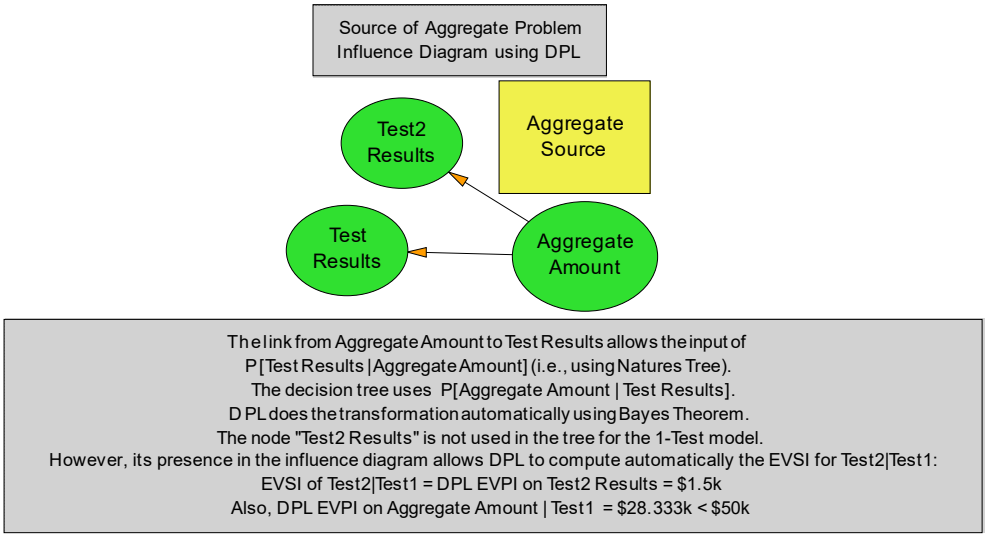
DPL Example 1: Basic Aggregate Source problem. No "test branch" in decision tree.

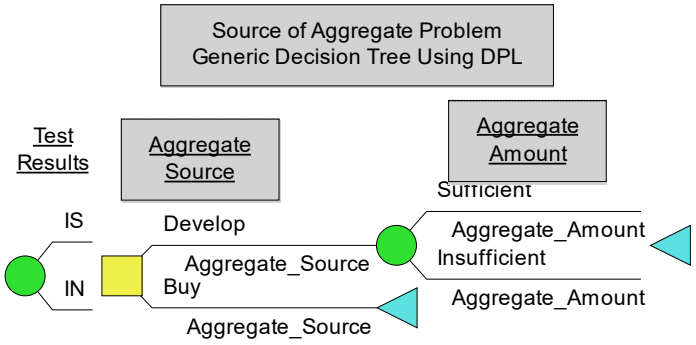


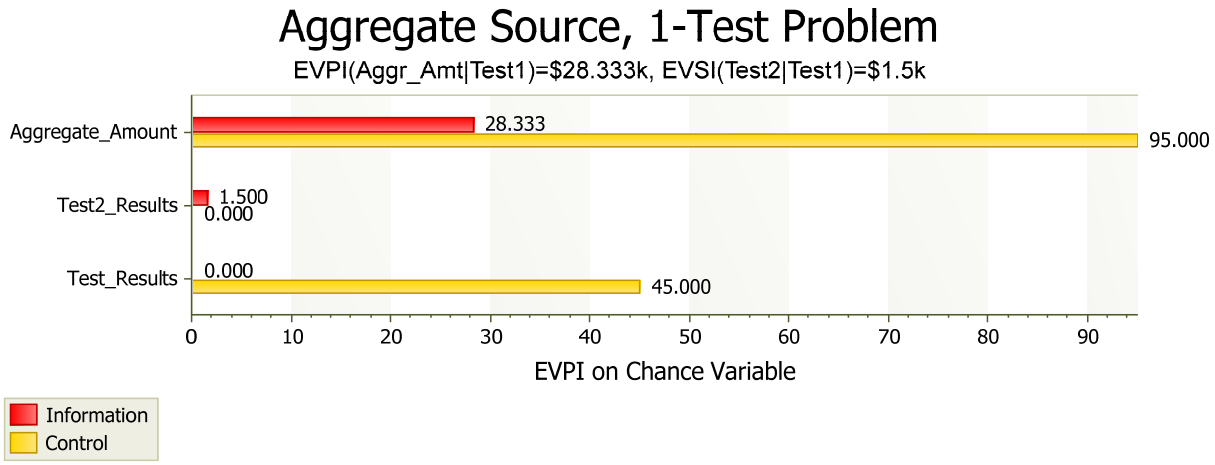




DPL Example 2: Aggregate Source problem with "Test Results" (Test1) in decision tree.







Sensitivity and Specificity for Binary Distinction Tests

Many decision models require *chance variables* or *random variables* that have only two outcomes. Such chance or random variables are called *simple* or *binary distinctions.* Similarly, *tests* or *experiments* to obtain information about the probabilities of the outcomes of random variables can also be binary or simple distinctions and have only two outcomes.

Consider a *binary distinction test* (a detector) that has *two test outcomes* for a *binary distinction* that also has *two outcomes*. In this case, the likelihood matrix (conditional probabilities) for the test (*detector*) requires the assessment of only two probabilities. In medicine, these are called:

$$\textit{Sensitivity} \text{ of the detector} = t = P[\text{IVirus} \mid \text{Virus}] = P[\text{IRainy} \mid \text{Rainy}]$$

$$\textit{Specificity} \text{ of the detector} = f = P[\text{INoVirus} \mid \text{NoVirus}] = P[\text{ISunny} \mid \text{Sunny}]$$

In general, the *value* of a test (detector) is its *value-in-use* to a decision maker and stems from the ability of the *information* provided by the test (detector) to change the best alternative for *the decision at hand*. Clearly, this ability and the value of the information increase as the test becomes more *accurate*, i.e., when the *sensitivity* and *specificity* of the test (detector) increase. When *sensitivity* and *specificity* reach 100% the test becomes *perfect* and its value reaches a maximum. The value of a test is specific to a decision situation. If none of the test outcomes can change which decision alternative is best for a given decision problem, then the test has no value.

Value of a Test and Material Information

Information is *material* if it may change the best decision alternative from that chosen without the information. *Material* information has value. *Immaterial* information has no value because it will not change the decision choice that would have been made anyway.

To have value, an *experiment* (a detector) must provide *material* information. An experiment cannot have value unless it is *material*; i.e., at least one of the experimental observations must have the possibility to change the alternative that would have been chosen by the decision maker.

Principles for Experimentation or Information Gathering

For experimentation or information gathering to be useful it must satisfy the following three requirements:

1. The results of information gathering must be ***relevant*** to distinctions of interest in the decision problem. *Information is relevant* if it changes the *conditional probabilities* of the distinctions of interest. It would not be prudent to spend money or time on a geologic exploration in Alaska, for example, if the decision is whether to develop a piece of land in Michigan as an investment.

2. The information gathering process must be ***material*** to the decision. The observations from an information gathering process must have the possibility of changing the decision that is made without those observations.

3. Information gathering should make *economic* sense in terms of the overall decision. I.e., its EVSI (value-in-use) must exceed the cost for obtaining it (value-in-exchange = market cost =cost of the test = CT).

Important: It is possible for information about one distinction to be *relevant* and *material* to a decision when considered alone; but when combined with information gathered about another distinction, to have no value whatsoever.

Example: The "Prefab Plant" decision problem

Perfect information about the "Demand in the first 2 years (D2)" has
EVPI(D2)= $19,963 when considered alone.

However, when combined with perfect information about the "Demand in the next 8 years (D8)", then perfect information about the "Demand in the first 2 years (D2)" has no value,
EVPI(D2|D8)= $0.

VALUATION
Market Prices, Value-in-Exchange, Value-in-Use
Personal Indifferent Selling / Buying Prices: PISP, PIBP

Market's Buying and Selling Prices

The *market's buying price* is what the market is willing to pay to buy something that a person owns. The *market's selling price* is what a person would have to pay in the market for a particular good or asset.

Note that the market's buying and selling prices are easily observable only for goods for which a market actually exists. The stock exchange is an example of a well-functioning market. The market's buying and selling prices for unique goods or assets (e.g., a particular house or a piece of art) may be observed only at the time of an actual transaction.

Value-in-Exchange

Value-in exchange is the market value of any good or service. The market value is specified by market buying and selling prices and is determined by supply and demand.

Value-in-Use

The *value-in-use* of something (such as a bicycle belonging to a student) represents its value to a *specific person*. It reflects the particular *situation* of that person (e.g., a not-so-well-to-do student). Most importantly, it reflects the *personal use* to which that asset is being put by that person (e.g., to provide transportation to university classes and other nearby places). The value-in-use for the same bicycle would be different for another person, or for the same student if his level of wealth changes, or if his personal use changes (take the bus to classes). An asset does not need to be already owned by a person to have value-in-use for that person. Something that a person needs and *wishes to buy for personal use* also has value-in-use.

Given the right circumstances, the value-in-use for something can significantly exceed its value-in-exchange. This is the point made by the well-known phrase from Shakespeare's Richard III, "*A horse*! *A horse*! *My kingdom for a horse*!" King Richard's exclamation that a horse is as valuable to him as his entire kingdom makes clear his specific personal circumstances *at that point in time*: he is in battle, he has lost his own horse, and he is afraid that he may lose his kingdom or his life if he remains on foot. He needs a horse now!

Because the concept of value-in-use reflects *personal use*, it applies only to those things that people use in their daily lives. Value-in-use *does not exist* for those things that a person buys and sells in the course of their business. For example, the bicycles bought and sold by a bicycle dealer do not have value-in-use for that dealer. This is because the dealer does not use those bicycles as a means of transportation in his daily life.

Two particular types of "value-in-use" are quite important in decision analysis and are defined below: The *Personal Indifferent Selling Price* (PISP) and the *Personal Indifferent Buying Price* (PIBP).

Personal Indifferent Selling Price (PISP)

A selling price arises when a person sells something that the person already owns (e.g., a bicycle) in exchange for money (cash). The word "indifferent" means that PISP is that special selling price at which the person is indifferent between

(a) Continuing to own that asset (e.g., the bicycle) and
(b) Giving it up in return for receiving the selling price (cash).

For example, a student might own a bicycle that he uses for transportation to school each day. The student's PISP for that specific bicycle is that price that would make the student *indifferent* between (a) keeping the bicycle and (b) selling it to someone else in return for money. The student's PISP for the bicycle is his *value-in-use* for the bicycle.

Desirable things like a useful bicycle have a positive selling price. Sometimes, however, people own undesirable things, such as trash or toxic waste that need to be disposed properly. The PISP for undesirable things can be a negative selling price. This means that the owner of an undesirable thing is willing *to pay* money to sell the item (and thus transfer its ownership and responsibility) to someone else.

The concept of PISP can be extended to all sorts of things, some of which are not as tangible as a bicycle. One such undesirable thing that most people own is *risk*. For example, the owner of a car also owns the risk of causing bodily injury or property damage to others and to himself, as well as the risk of theft or damage to the car. People sell these risks to insurance companies for a negative selling price. A car owner pays an insurance premium to an insurance company to transfer the consequences of car–related risks from himself to the company.

A person's PISP for something can be defined as follows:

- The PISP (selling price) for a desirable asset is the *least* amount of money the owner would be willing to accept in order to forgo (give up) ownership of the asset.
- The PISP (selling price) for an undesirable asset (or a liability) is the *most* amount of money a person would be willing to pay to transfer ownership of the asset (or liability) to someone else.

Notice that the above two definitions for PISP (selling price) are actually the same if we define a selling price algebraically so that *cash received* is *positive* and *cash paid out* is *negative*.

- The PISP (selling price) of an asset is the *least* amount of money (in an algebraic sense) in return for which a person would be willing to forgo (give up) ownership of the asset.

For a desirable asset, the above definition is self explanatory. For an undesirable asset, the algebraic value of "the least amount of money" is a negative number and represents "the *most* the owner would be willing to pay" to get rid of the asset.

Meaning of the terms in PISP — Personal Indifferent Selling Price

- It is a "*selling price*" because the seller is going to exchange ownership of something he owns for money.
- "*Indifferent*" means that the seller would not accept any lower amount of money as a price; i.e., at the PISP price he really does not care *whether he sells or not*.
- "*Personal*" means that the seller is forgoing (giving up) *personal use* of the asset. For a car, this would be transportation and possibly enjoyment. The seller is not thinking of himself as a broker (e.g., a car dealer) or considering any other possible selling opportunities for the asset. The key concept here is that of personal use of the asset. The word "*personal*" also emphasizes that PISP depends entirely on the asset's owner. Some other person with exactly the same asset may have a totally different PISP for it.

A person's PISP (selling price) for something *does not depend* on how much the market is willing to pay for it. I.e., it does not depend on the market's buying price for the asset. It only depends on the asset's "*value-in-use*" to its owner. Moreover, an owner's PISP (selling price) *does not depend* on what that person paid for it (e.g., it could have been a gift). PISP is simply the price that makes the owner indifferent between continuing to have something or not.

In summary, the term PISP does not involve any notion of a market price, gain, or any other price except that one particular price that makes a person indifferent between keeping and using something vs. selling it at that price. It is important to stress this point and differentiate it from a *corporate indifferent selling price*, CISP, that may be the result of a *decision process*. As we will see later, a company considering the sale of a drug patent, for example, should determine its CISP through careful analysis.

Example: PISP, Value-in-Use and Value-in-Exchange

A situation that illustrates clearly the difference between "value-in-use" and value in exchange" is that of the amateur pilot of a biplane who also happens to be a steel-cable merchant.

As an extreme example, consider the pilot's PISP (selling price) for the steel cable that connects the steering lever of his biplane to the rudder *while the plane is flying at 120 mph at 1000 feet above the ground*. It is obvious that while in the air, the pilot's minimum selling price for that cable would be astronomical and would reflect the fact that at that moment he is actually *using* the steel cable to control the plane and avoid a crash. It represents his "*value-in-use*" for the steel cable "while he is actually *using it to fly*". Obviously, an overwhelming consideration for setting this PISP (selling price) is the possibility of death or at least serious injury.

On the other hand, if the same person was on the ground and in his shop as a steel-cable merchant and someone wanted to buy the identical length of steel cable (cut from a spool of brand new steel cable), the price would be quite low and competitive with other similar shops. When the pilot is acting as a merchant in his shop, the selling price would be its value-in-exchange and would be determined by the prevailing market prices for steel cable (supply and demand). The market selling price will have nothing to do with his *personal use* of the cable.

Example: A Dramatic Change in PISP (Value-in-Use) Over Time

The biplane is now in its hangar undergoing routine maintenance. Clearly, the pilot's PISP for the same steel control cable has changed. When the plane was in the air, the PISP was very high. Now that the place is on the ground, the PISP for the steel control cable can even be negative.

As part of maintenance, for example, it is common to replace the steel control cable with a new one. In this case, the old cable is trash and the pilot might be willing to pay a small amount to his mechanic to get rid of it. Technically, the pilot is selling the old steel cable to his mechanic for a negative selling price (he pays the mechanic money to dispose of the cable). Similarly, when replacing tires, a car owner is willing to pay the shop a fee for disposing of the old tires.

Personal Indifferent Buying Price (PIBP)

A buying price arises when a person buys something that he does not already own. The word "indifferent" means that PIBP is that special buying price at which a person is *truly indifferent* between:

(a) Buying that thing in exchange for that amount money (the PIBP), and

(b) The status quo: I.e., the person's current state of not owning that thing and not paying any money (i.e., the person acquires nothing and pays no money, there is zero change in assets).

For example, a student might be considering buying a particular bicycle that he will use for transportation to school each day. The student's PIBP for that bicycle is that special buying price that would make the student *indifferent* between (a) paying to money to obtain that specific bicycle and (b) walking away without the bicycle and with all his money intact. The student's PIBP for the bicycle is his *value-in-use* for the bicycle.

Of course, not all assets that a person can buy are necessarily desirable. Some could be undesirable or represent liabilities or obligations that the person did not originally own.

As a result, a person's PIBP for something he does not already own can be defined as follows:

- The PIBP (buying price) for a desirable asset (such as a bicycle) represents the *most* money he would be willing to pay to acquire ownership of that asset.
- The PIBP (buying price) for an undesirable asset (such as someone else's trash) represents the *least* amount of money *he would accept as payment* to take on ownership (or responsibility) for that undesirable asset (or obligation, or liability).

Notice that if we define *cash paid out* as a *positive* buying price and *cash received* as a *negative* buying price, then the above two definitions for PIBP (buying price) for desirable and undesirable assets (obligations) are actually the *same*.

> The PIBP (buying price) of an asset is the *most* amount of money (in an algebraic sense) that a person would be willing to pay to acquire ownership of the asset.

For a desirable asset, the above definition is self-explanatory. For an undesirable asset, the algebraic value of "the *most* amount of money that a person would be willing to pay" is a

negative number that represents "the *least* amount of money that the person would accept as payment" to take ownership of the undesirable asset (liability, obligation). I.e., in this case the buying price is "negative" and represents money received and not money paid out.

Just like for PISP, it is equally important to establish the exact meaning of all the words in PIBP (personal indifferent buying price):

- It is a "*buying price*" because the buyer is going to exchange money to acquire ownership of something he does not already own.
- "*Indifferent*" means that the buyer would not pay a greater amount money as a price; i.e., he truly does not care whether he buys or not at the PIBP price.
- "*Personal*" means that the buyer is considering the acquisition of an asset for *personal use*. For a car, this would be transportation and possibly enjoyment. The buyer is not thinking of himself as a broker (e.g., a car dealer) or considering any other possible selling opportunities for the asset. The key concept here is that of personal use of the asset. The word "*personal*" also emphasizes that PIBP depends entirely on the buyer of the asset. Some other buyer can have a totally different PIBP for the same asset.

A person's PIBP (buying price) for something *does not depend* on its selling price in the market. It only depends on its "*value-in-use*" to the buyer. Moreover, a buyer's PIBP (buying price) for something *does not depend on what he can sell it for* or what he has to actually pay for it in the market. PIBP is the price that makes the buyer indifferent to becoming an owner or not.

The previous examples of trash, toxic waste, or an old car rusting in the owner's driveway can also be considered from the point of view of someone who does not currently own them, but will take on responsibility for their disposal. Someone's rusty old car would have a negative PIBP for the person who will tow and dispose of it. I.e., this person would demand (rather than pay) money to take on the car's ownership and responsibility. In this case, the original car owner has to pay to get rid of the car (i.e., the owner sells at negative price) and the new owner is actually getting paid to tow the car away (i.e., the buyer is buying the car at a negative price). A similar situation exists for persons who take on the responsibility to dispose trash or toxic waste.

PISP and PIBP *change over time* as the value of *personal use* of the asset in one's daily life changes. The exclamation by Richard III, "*A horse*! *A horse*! *My kingdom for a horse*!" shows that *at that point in the battle* his PIBP for a horse was *at least* equal to his entire kingdom! If Richard had survived the battle (he did not), his PIBP for a horse would have been much less.

Transactions—Buying and Selling

In theory, people have personal indifferent selling prices (PISP) for everything they own and personal indifferent buying prices (PIBP) for everything that they do not own but might buy.

When should a desirable asset be purchased?
When its PIBP ("value-in-use") is more than its "value-in-exchange" (market price).
When should a desirable asset be sold?
When its PISP ("value-in-use") is less than its "value-in-exchange" (market price).

Of course, people's needs, tastes, and assets change and so do opportunities, information, and circumstances that affect wealth. As a result, people's personal indifferent buying and selling prices (PIBP, PISP) change over time, too. And such changes happen often and quickly.

PISP and PIBP in Construction

The PISP and PIBP concepts presented above can be extended to all goods and services, including constructed facilities, design services and construction services.

We are used to thinking that contractors sell construction work and owners buy construction work. Strictly speaking, however, this is only true for facilities that are constructed and paid for by contractors with the intention of selling them to others. A typical example is that of a house that a contractor has already built in speculation and which is now being sold to a prospective homeowner. In this case, the contractor is selling a finished product that he constructed with the intention of selling it. The situation is similar to a dealer selling a car. The contractor is selling a house he owns and the future homeowner is buying a finished home.

In the more common case, however, where the owner and his designer obtain permits and financing, and produce bidding documents to solicit proposals (bids) for construction, the roles of buyer and seller are actually reversed. At bidding time, the future *owner* of the facility is actually *selling* the obligation to perform construction according to plans and specs and the *contractor* is *buying* the obligation to perform construction according to contract documents.

Clearly, a contractor's PIBP is the *least* amount of money he would accept as *payment* to take ownership of the obligation to perform construction work that he currently does not have to do. In this case, the contractor does not originally own the obligation to perform that construction. The owner has this need. When a contractor signs a construction contract for a given price he is *buying* from the owner the obligation to do very specific work at a *negative* buying price where the owner pays and the contractor receives money. This money is a *negative buying price* for the contractor and also a *negative selling price* for the owner.

To make it clear that at bidding time the contractor is actually making an offer to *buy* "the contractual obligation to perform" at a *negative* buying price, consider his alternatives: (a) to buy the obligation to perform construction work and get paid for it or (b) to walk away (no obligation to do work, no payment for the work, *no change in the contractor's asset position*). The specific project is not the buyer's (the contractor's) problem. It is the seller's (the owner's) problem.

For the owner, the PISP for the work is the most he would be willing to pay to sign a construction contract and transfer the need/obligation to perform construction work to a contractor. In absolute value terms, if the owner's PISP is larger than at least one contractor's PIBP then there is the possibility that an agreement may be reached. Otherwise, if the owner's PISP is below the PIBP of every contractor then no agreement is possible.

Exactly the same situation exists between a *general contractor* and a *subcontractor* when they are about to sign *a subcontract*. In this case, *the general contractor is a seller* of work to do (like the owner above) and *the subcontractor is a buyer* (like the contractor above). In a subcontract, the sub buys the obligation to do part of the GC's work obligation (at a negative buying price).

Cycle of Ownership — PIBP = PISP

A cycle of ownership is a mental exercise where a person is involved in the instantaneous buying and selling of an asset at exactly his PIBP and PISP for that item (no more, no less). The purpose of this mental exercise is to illustrate that:

Around a cycle of ownership, PIBP = PISP = value-in-use

This is an extremely important result. At any given point in time, PIBP = PISP.

Values around a Cycle of Ownership

One proof for the equality of PIBP and PISP around a cycle of ownership starts like this:

[Particular Asset] + [Rest of assets] ~ [PISP] + [Rest of assets]

[House] + [Rest of assets] ~ [PISP for the house] + [Rest of assets]

Starting on the left, a person owns a particular asset (e.g., a house) plus the rest of his assets. If the person sells the house at a price *exactly* equal to his PISP, his situation becomes the one shown on the right. The symbol "~" indicates that, by the very definition of PISP, the person is *truly indifferent* between his "before" state (the left side) and his "after" state (the right side).

At this point, the person is on the right side of the indifference statement. Let us assume that this person will *immediately* buy back the asset he just sold (the house) without any other changes (i.e., no changes in wealth, personal needs, information, etc.).

Since the person is *truly indifferent* between the two sides, he would be indifferent between the status quo (the "before" state shown on the right side) and paying PISP to buy the house and move to the "after" state (the left side). So, PISP is the buying price for which the person is indifferent between buying the house or not. However, by definition, the special price for which this indifference exists is PIBP, so

[PISP] = [PIBP]

Given the equality of PISP and PIBP around a cycle of ownership, we can also write the above indifference statements as:

[Particular Asset] + [Rest of assets] ~ [PIBP] + [Rest of assets]

[House] + [Rest of assets] ~ [PIBP for the house] + [Rest of assets]

Important: for PISP and PIBP to be equal, a cycle of ownership must not change a person's value system. This means that there should be no passage of time, no new information, no new needs, and no change in the person's state of wealth. In other words, a cycle of ownership exists only *ceteris paribus*, i.e., when all other things, including time, are equal or held constant.

A cycle of ownership can only occur when the price paid to buy an item is *exactly* equal to its PIBP and when an item is sold at a price *exactly* equal to the owner's PISP. If the price paid or received is more or less there cannot be a cycle of ownership because there is a *change* in the person's *wealth*. When a person becomes richer or poorer there is no reason to believe that his valuation for things remains the same. In fact, experiential evidence points to the contrary.

Similarly, a cycle of ownership does not exist if a person receives new information or there is a passage of time (which can change the use, needs or circumstances). Such changes can alter a person's value-in-use for a particular asset.

Effect of a Person's Wealth on Values and the Cycle of Ownership

A person typically values things differently as his wealth (or asset position) improves. For example, a person who used to own one suit would have a totally different selling price for the same suit if he now owns twenty suits. It would not be surprising to find out, for example, that that person's PISP for the original suit has decreased considerably. This can explain in part why people donate their "old" clothes (some of which they may have never worn).

A similar phenomenon would involve a student and his bicycle (his primary means of transportation while at school). A week after the student graduates, he earns considerably more and buys a car. His PISP (selling price) for the "old" bicycle may now be considerably less.

Let us now consider a college professor whose PIBP for a brand-new Bentley (a luxury British car) is $50,000. As the market price for a new Bentley is $100,000, the professor does not own one. (Remember that the professor is not a car dealer. He would buy the Bentley to use it.) The professor's PIBP for such an expensive car is so low because he feels that he really cannot afford the maintenance and repair of a Bentley nor the gasoline he would need to buy to use it for transportation. Moreover, he is also concerned that someone may scratch it in the parking lot or just steal it. Insurance costs would be very high. He is also worried that his children will spill drinks and food everywhere. Family trips in such an expensive car may prove to be challenging.

Suppose that the professor's research has led to a successful invention that has made him a multimillionaire. The professor may now feel that the time to buy the Bentley has finally arrived. This means that given his new wealth, his PIBP may now exceed the $100,000 sticker price and he can buy the car.

In both cases, we can make the same observation: A change in the [student's | professor's] wealth has changed their personal indifferent [selling | buying] price for the [bicycle | Bentley]. We cannot always predict which way PISP and PIBP will move as wealth increases. They can go either way: increase or decrease.

Effect of a Person's Wealth on the Cycle of Ownership

An interesting question is whether we can construct a cycle of ownership around the professor's original PIBP of $50,000. What makes this a challenging case is that this PIBP is not high enough to actually buy a $100,000 Bentley. Such a cycle of ownership would look like this:

[PIBP=$50,000] + [Rest of assets] ~ [Bentley*] + [Rest of assets]

There is such a discrepancy between the PIBP and the market price of the Bentley that extra care is required to understand what the above indifference statement really means. It is clear from the professor's statement that he is indifferent between his "before" state on the left, where he has $50,000 in the bank and the "after" state on the right where he has bought the Bentley for

personal use (in exchange for his PIBP of $50,000). Notice that the Bentley on the right is shown with a (*) next to it to indicate that this is not just any car. It is a car that the professor must keep and use in his daily life. He cannot turn around and sell it. He has to live with it day-in-and-day-out and has to pay all its costs for insurance, gasoline, maintenance and repairs.

Most people find it difficult to understand this situation because the $100,000 market value of the Bentley on the right is so high. How can $50,000 on the left side be equivalent to something worth $100,000 on the right side?

This is indeed a valid concern. To make the situation more realistic and easier to comprehend, let us assume that for advertising purposes, the Bentley dealer is willing to subsidize the professor's purchase of the Bentley and accept his offer of $50,000 for the car. The catch is that (a) the professor cannot sell the car to anyone else other than the dealer, and (b) he must use it in his daily life so people can see it (i.e., he has to drive it like a regular car). These provisions make the (*) next to the Bentley easier to understand. They also make much clearer the *personal use* conditions that the professor really had in mind when he stated his PIBP.

The professor may not be willing to pay more than $50,000 for the $100,000 Bentley, but this is because he has to keep it, maintain it, and use it for a certain number of years. He cannot just buy and sell it to make a profit. Given the subsidy arrangement with the Bentley dealer, it is easy to see that the above indifference statement can now work equally well from right to left. Immediately after the professor buys the subsidized Bentley, the dealer offers him the opportunity to change his mind and sell the car back to the dealer. The professor's PISP (selling price) would be exactly equal to his PIBP of $50,000. He would be indifferent between a life with the expensive gas-guzzling Bentley and getting his money back.

Suppose that while the professor is considering whether to sell the Bentley back to the dealer, the Bentley factory representative hears this unusual story. He immediately appreciates the advertising potential of the situation and offers the professor an even better deal. He can have his $50,000 back and also drive the car for as long as he wants! Moreover, when he is tired of it he can sell it back to the dealer. He can even sell it back now if he wants!

What has really happened? The professor has just become the owner of the Bentley for no money, i.e., for a price that is definitely less than his PIBP. As a result, the original cycle of ownership (shown above) no longer applies. Moreover, the Bentley is worth something, so the professor's wealth has increased. Because of this change in wealth, his PISP and PIBP for the car (and everything else) have also changed. (Notice that the professor's wealth has not necessarily increased by $100,000 because he can only sell the car back to the dealer who may not be willing to buy it at the full price.)

The interesting question is this: If the professor was to sell the car back to the dealer immediately, what would his PISP for the car be now? We do not know the answer to that question. Chances are it is more than $50,000. It could be $60,000, $70,000, or whatever value would make the professor indifferent between keeping the car (and using it) and selling it back to the dealer. This new PISP (like the old PISP) is personal and must be assessed by the professor. It has nothing to do with the dealer or the market price for the car.

Effect of New Information on Values and the Cycle of Ownership

Continuing the example of the professor and the Bentley, suppose that right after the professor has stated that his PIBP for the car is $50,000, he listens to the news and finds out that gasoline prices are expected to double for the foreseeable future. This new information has a significant bearing on the operating costs of a gas-guzzler like the Bentley. It would not be surprising if this information reduces the professor's PIBP for the car below $50,000.

Let us now assume that this information arrived just *after* the professor bought the car at the dealer's subsidized price of $50,000. The arrival of this information has invalidated the original cycle of ownership because the PIBP for the car is now less than $50,000. This also means that we should no longer expect the PISP for the car to remain at $50,000. Given this new information, we can construct a new cycle of ownership but with lower values for PISP and PIBP. The important thing is that PISP will again equal PIBP around this new cycle of ownership. For the professor, the new information means that *his wealth has decreased*.

Extensions

The PIBP and PISP concepts presented above can be extended to clarify many issues of valuation and to allow the analysis of situations that at first may appear challenging.

As an example, let us introduce some more twists to the offer made to the professor by the Bentley dealer. In this new case, the professor is provided with the opportunity to buy the Bentley for money but he must use it in his daily life for exactly ten years (for advertising). After that he is free to sell the car to anyone he wants (not just the dealer).

The professor's PIBP now consists of two parts. His PIBP for ten years of use (which includes the enjoyment of luxury transportation services, but also maintenance costs, operating costs, insurance costs, etc.) and what he believes the present value to be today for a used Bentley sold ten years from now. The first part is subjective. It is the value to the professor of enjoying the Bentley for exactly ten years but also being responsible for its upkeep.

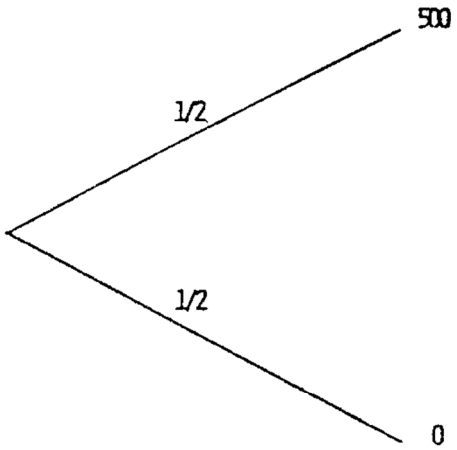
In contrast, the second part can be quite objective and could be evaluated easily today. For example, if there was a market in "*used Bentley futures*", it would allow the professor to sell today "*the ownership of a used Bentley ten years from now*". In the absence of such a market, the professor could try and get a ten-year loan and use the Bentley as collateral. That should provide him with a good indication of what the bank thinks a ten-year-old used Bentley is worth today. Finally, the professor could consult specialized catalogs that track the prices of used cars by year at the wholesale and retail markets and do his own present-value calculations.

Expected Value of Information = PIBP for that Information

The *Expected Value of Perfect Information* (*EVPI*) about the future outcome of a chance node is its PIBP. Similarly, the *Expected Value of Sampled Information* (*EVSI*) provided by a test is its PIBP. The *EVPI* and the *EVSI* are the *most* the decision maker would be willing to *pay* to *buy* that information (before it becomes known to him). I.e., the *EVPI* and the *EVSI* are the buying prices ($EVPI \geq EVSI$) that make him *indifferent* between buying the information or not.

© 2024 *Photios G. Ioannou*

Utility Theory



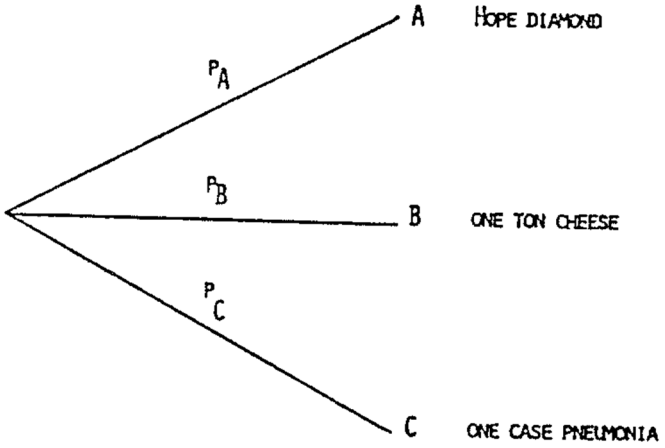
A COIN-TOSSING LOTTERY.

Shorthand notations:

($500, 0.5; $0)

or

(0.5, $500; 0.5, $0)



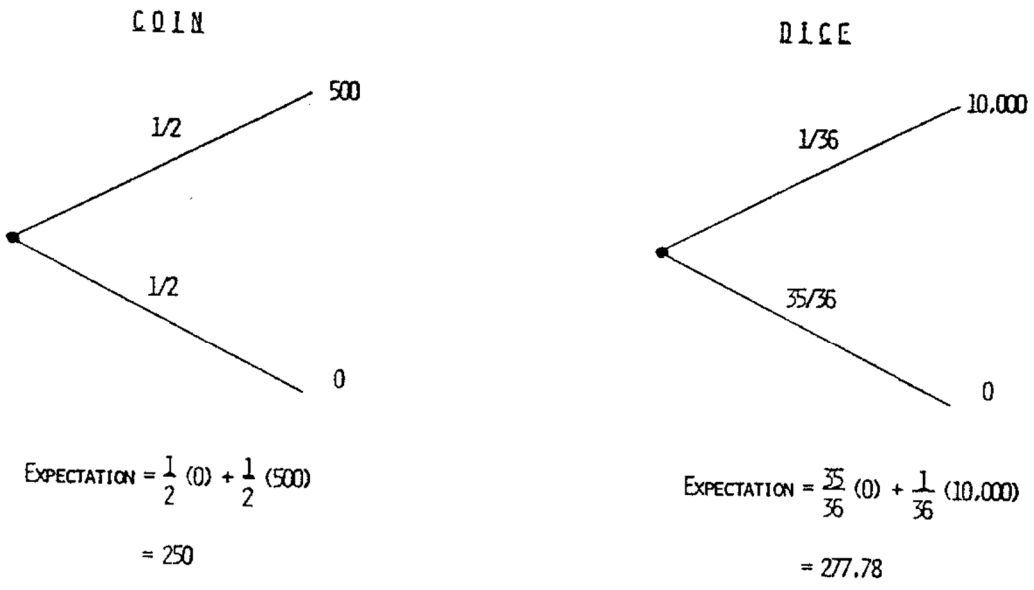
AN UNUSUAL LOTTERY.

Shorthand notations:

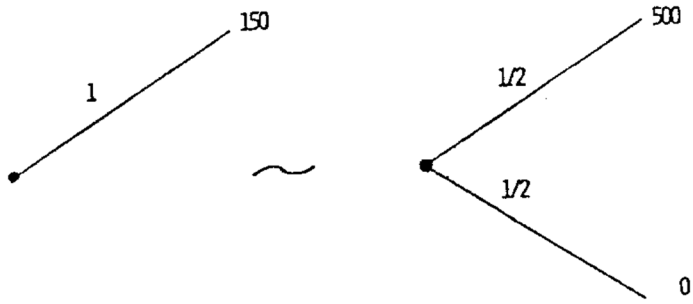
(A, p_A; B, p_B; C)

or

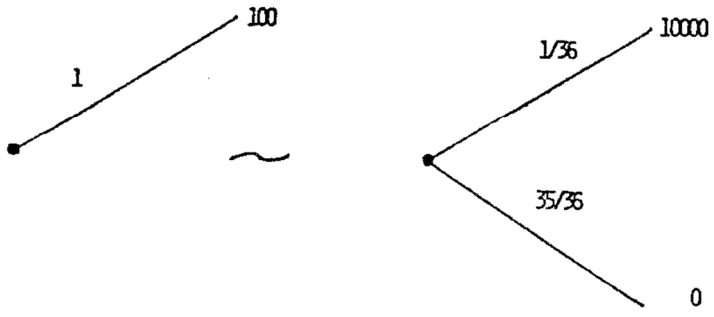
(p_A, A; p_B, B; p_C, C)



COMPUTATION OF EXPECTATION.



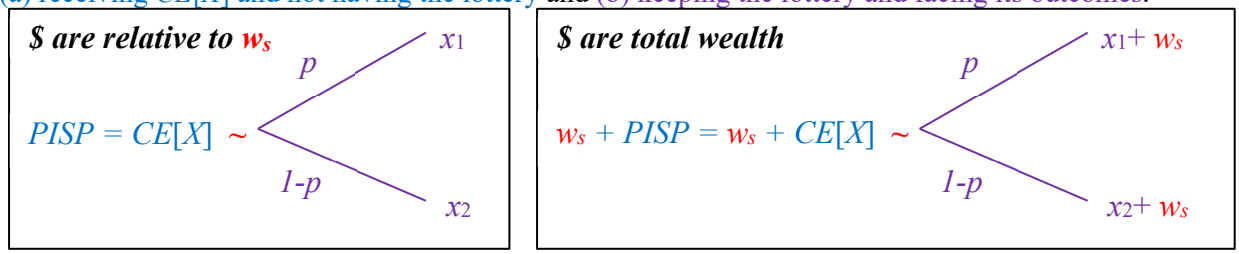
CERTAIN EQUIVALENT OF THE COIN LOTTERY.



CERTAIN EQUIVALENT OF THE DICE LOTTERY.

Certain Equivalent = PISP (Personal Indifferent Selling Price) for a lottery you own

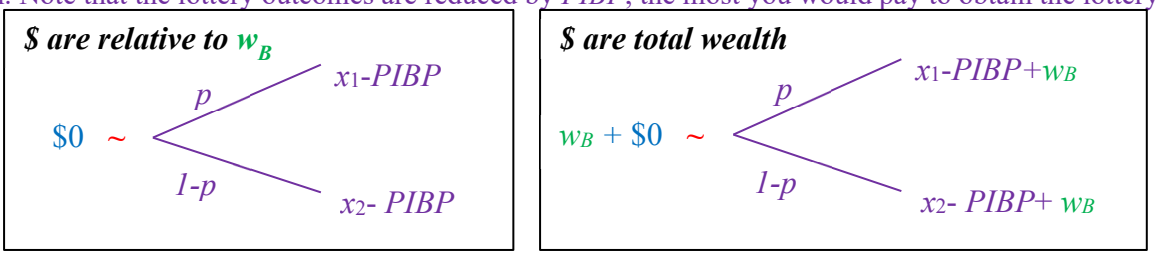
The **Certain Equivalent**, $CE[X]$, for a lottery you already own (when the rest of your current wealth is w_S), equals your **PISP**. It is the amount of money that would make *you* indifferent (~) between (a) receiving $CE[X]$ and not having the lottery and (b) keeping the lottery and facing its outcomes.



"Receiving the certain equivalent" has algebraic meaning. A positive certain equivalent is the minimum amount of money you would accept in order to give up (sell) a desirable lottery that you already own. A negative certain equivalent is the maximum amount of money that you would be willing to pay to avoid an undesirable lottery that you currently must face. In this case, the selling price is negative because you are willing to pay someone to take away this undesirable lottery and remove the associated risk of losing money. For example, a contractor that has to perform risky work (an undesirable lottery with cost (negative) outcomes) would be willing to pay a sub up to $CE[X]$ to do the work so that the contractor will not have to do that work and face its uncertain cost outcomes.

PIBP (Personal Indifferent Buying Price) for a lottery you do not already own

Let your current wealth be w_B. The **PIBP** for a lottery that you do not currently own is the most *you* would be willing to pay to acquire that lottery. I.e., you would be indifferent between (a) staying at your current wealth w_B (i.e., \$0 net change in wealth) and (b) acquiring that lottery that you do not currently own. Note that the lottery outcomes are reduced by *PIBP*, the most you would pay to obtain the lottery.



The *PIBP* also has algebraic meaning. A positive *PIBP*, is the most you would *pay* to buy a desirable lottery that you do not currently own. A negative *PIBP* is the minimum you would accept *to be paid* to obtain an undesirable lottery that you do not currently have to face. For example, a subcontractor needs to be paid at least *his PIBP* to sign a contract for work with uncertain costs (an undesirable lottery). In this case, the lottery outcomes (x_i –*PIBP*) represent the subcontractor's possible profit outcomes, v_i, because $x_i < 0$ are costs, and (–*PIBP*) > 0 is revenue, making profit outcomes positive: $v_i = (x_i - PIBP) > 0$.
Important: In general, ***PISP* ≠ *PIBP***. The *PISP* for a lottery you already own does not equal the *PIBP* for the same lottery if you do not own it, because your **current wealth** is different ($w_S + PISP \neq w_B$).
Important: In a *Cycle of Ownership*, we have *PISP* = *PIBP*, because $w_S + PISP = w_B$, and $w_B - PIBP = w_S$.
When $w_B - PIBP = w_S$ ➔ *PIBP* lottery = *PISP* lottery ➔ $w_B + \$0 = w_S + PISP$ ➔ $w_S + PIBP = w_S + PISP$ ➔ $PIBP = PISP$.

Risk Premium

The **Risk Premium**, $RP[X]$, of a lottery is that portion of expected value you are willing to forego (i.e., to give up) when you trade a lottery for its certain equivalent (to avoid the risk of the lottery). For a risk-averse person the risk premium is positive, while for a gambler it is negative.

$$RP[X] = E[X] - CE[X]$$

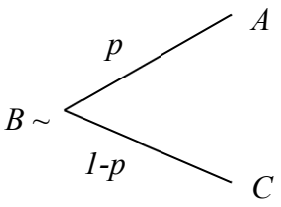
Utility Theory Axioms

1. **Orderability Axiom**
   The decision maker should be able to express a preference between the lottery prizes A and B. He must be able to state one of the following:

   $$A \succ B,\ A \sim B,\ A \prec B$$

   Furthermore, preferences should be transitive: If $A \succ B$ and $B \succ C$ then $A \succ C$.

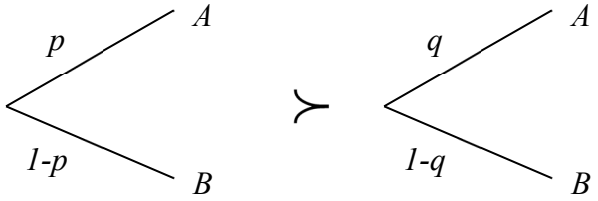
2. **Continuity Axiom**
   If $A \succ B \succ C$, then there must exist some probability p such that:



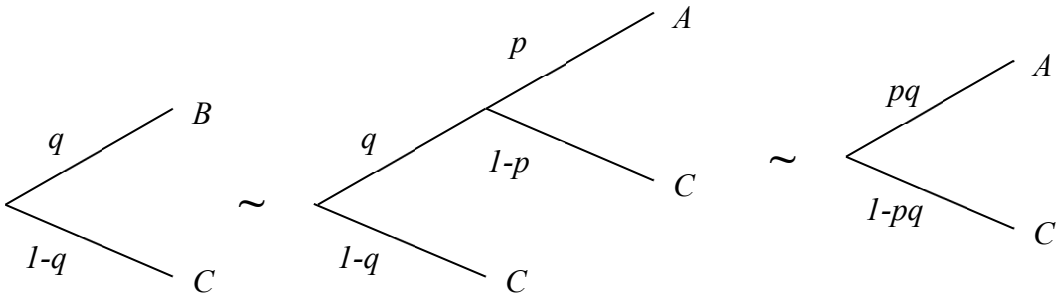
Note: The probability p is called a "*preference probability*". Later, it will be shown that p is very close to the concept of *the utility of outcome B*, $u(B)$. Outcome B is called the *Certain Equivalent* of the lottery. The *orderability axiom* describes preferences in a world without uncertainty. The *continuity axiom* describes preferences in a world with uncertainty.

3. **Substitutability Axiom**
A decision maker must be truly indifferent between the Certain Equivalent of the lottery B and the lottery itself $(p, A; 1\text{-}p, C)$. The lottery and its certain equivalent must be *interchangeable* without affecting preferences. *Substitutability* is a "*Do you really mean it*?" axiom.

4. **Monotonicity Axiom**
   If $A \succ B$, then the following is true only if and only if $p > q$:



5. **Decomposability Axiom**
   A compound lottery can be turned into an equivalent simple lottery using the laws of probability theory. This is a "*No fun in gambling*" axiom.



Utility Curve

A person whose preferences satisfy the five utility theory axioms can encode their preferences under uncertainty in a *utility curve* that assigns a utility value $u(x_i)$ to every prize x_i.

- The utility $u(L)$ of a lottery L equals the expected utility of its prizes x_i.

$$u(L) = \sum p_i u(x_i)$$

- The decision maker prefers the lottery with the higher utility.

$$L_1 \succ L_2 \quad \Leftrightarrow \quad u(L_1) > u(L_2) \quad \Leftrightarrow \quad \sum p_{1i} u(x_{1i}) > \sum p_{2i} u(x_{2i})$$

- The utility of a lottery $u(L)$ equals the utility of its certain equivalent $CE[X]$.

$$u(CE[X]) = u(L) = \sum p_i u(x_i)$$

- A utility function is like a "preference thermometer". The utility values have no meaning in themselves. They serve only to compare the desirability of lotteries.
- Like for a thermometer scale, we can multiply a utility curve by a positive number and add a constant to it without affecting the preferences it represents.

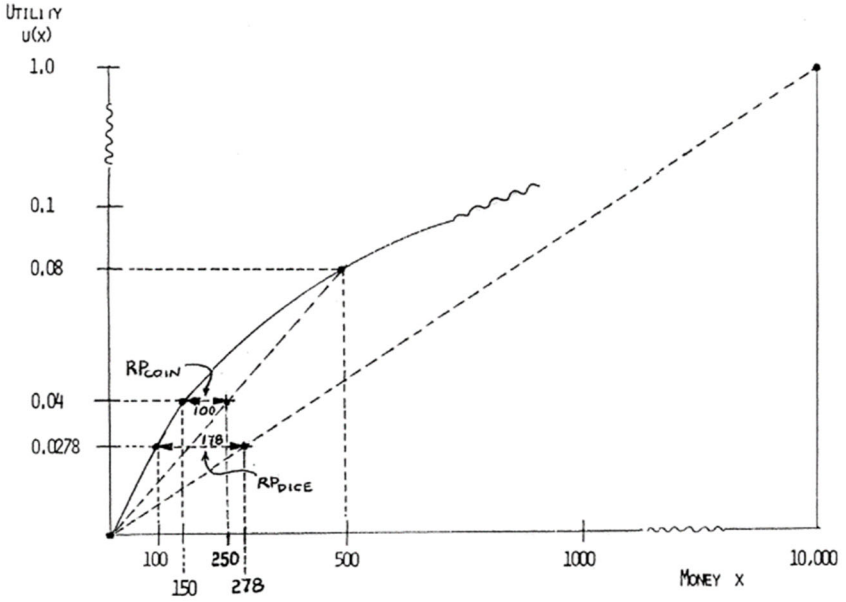
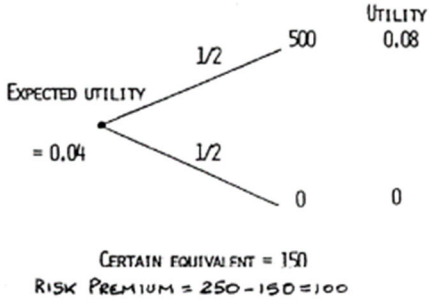


FIGURE 11 A UTILITY CURVE.

COIN



DICE

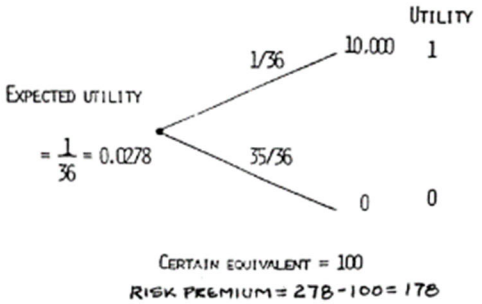


FIGURE 12 USING THE UTILITY CURVE TO CHOOSE BETWEEN LOTTERIES.

Linearity of Value

Up to now, the *measure of value* has been monetary value in dollars. The criterion for making decisions was to select the decision alternative with the highest *expected monetary value* (EMV).

Using EMV requires that the decision maker has a linear value system. A person's value system is linear if twice *any amount* is worth twice as much to him, that is, if he would be indifferent between a wager of any amount for a 50% probability of winning double that amount.

Thus, most people can use EMV only when:

(a) The monetary outcomes of all prospects are small relative to the decision maker's wealth.

(b) The monetary outcomes of all prospects are large, but the monetary differences between these possible outcomes are small relative to the decision maker's wealth.

Most people do not have a linear value system when making decisions that involve sums of money in the order of the decision maker's total assets. A typical person would not be willing to bet all that he owns for a 50-50 chance of ending up with (a) twice what he owns (or its equivalent monetary value) or (b) nothing.

Thus, most people should use a nonlinear value scale (a utility curve) when:

(c) Monetary outcomes of prospects are both small and large such that *the monetary differences between these outcomes are large* relative to the decision maker's wealth.

Thus, a person can use EMV to make decisions if (a) or (b) are true, but should encode and use his own personal utility curve to make decisions if (c) is true.

Utility

A person's *utility curve* represents the personal value that a person ascribes to money (or any other attribute) when faced with uncertain outcomes (gambles).

Each person has a different utility curve for money.

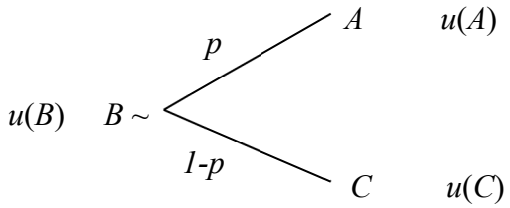
A decision maker's best alternative is the one with the highest expected utility (EUV) based on the decision maker's own utility curve.

Utility Curve Assessment

- To assess a person's utility curve, we use a *utility scale* that is similar to a temperature scale.
- To establish a utility scale for a person's utility curve
- We select two amounts of money *Hi* and *Lo* and we assign their utility values arbitrarily.
  - $u(Lo)=0$ and $u(Hi)=1$
  - This is like assigning 0ºC to the freezing of water and 100ºC to the boiling of water.
  - So, sometimes we assign $u(Lo)=0$ and $u(Hi)=100$.
- Typically, we select *Hi* and *Lo* such that $Lo \le x_i \le Hi$ for all outcomes x_i involved in a particular decision.
- This means that $0 \le u(x_i) \le 1$ for all outcomes x_i involved in the decision.
- The problem in utility assessment is:
  How to find the utility $u(x_i)$ for other values x_i in the interval from *Lo* to *Hi*.

Utility and the Certain Equivalent of a Lottery

Consider a situation that involves the three amounts of money A, B, C, where $A > B > C$ (i.e., A is preferred to B, and B is preferred to C). On one hand, we have a sure amount of money B. On the other hand, we have a lottery (a gamble) that pays off either A or C as prizes with probabilities $P[A] = p$ and $P[C] = 1\text{-}p$ respectively.



Suppose that a decision maker is indifferent between (a) receiving the sure amount B, and (b) playing the lottery with risky outcomes A and C., Then, for this person, B is the *certain equivalent* (or *personal indifferent selling price PISP*) for this lottery. This means that B is the least amount of money that the decision maker would accept for sure to sell this lottery to someone else (if the decision maker did indeed own this lottery).

Notice that this indifference statement involves a fixed amount of money B (a sure thing) and a lottery that has uncertain outcomes A and C (a risk). The indifference between the certain equivalent and the lottery can also be written more compactly as:

$$B \sim (A, p; C)$$

Where the symbol "~" indicates that the decision maker is indifferent between the certain equivalent B and the lottery $(A, p; C)$. This equivalence implies that the utility of B equals the expected utility of the lottery:

$$u(B) = p\, u(A) + (1\text{-}p)\, u(C)$$

Preference Probability and Utility

The above figure was also used when presenting the *Continuity Axiom*.

In that case, the probability p was called the *preference probability* and it was mentioned that the *preference probability p* is very close to the concept of the utility of outcome B, $u(B)$.

We can see this by setting $Hi{=}A$ and $Lo{=}C$, and assigning arbitrarily:

$$u(A){=}1 \text{ and } u(C){=}0$$

Then, from the last equation we have that:

$$u(B) = p\, u(A) + (1\text{-}p)\, u(C) = p = \textit{preference probability}$$

Thus, the points on a utility curve show the preference probabilities for outcomes (between *Hi* and *Lo*) that are the certain equivalents for simple lotteries with outcomes *Hi* and *Lo*. Thus, each point on the utility curve is a different (x =certain equivalent, $y{=}u(x){=}p$=preference probability).

Utility Assessment Procedures

To assess a person's utility function for a particular attribute (e.g., dollars profit), select two values *Lo* and *Hi* that bracket all outcomes x_i involved in the decision, i.e., $Lo \le x_i \le Hi$.

To these values assign arbitrarily, $u(Lo)=0$ and $u(Hi)=1$.

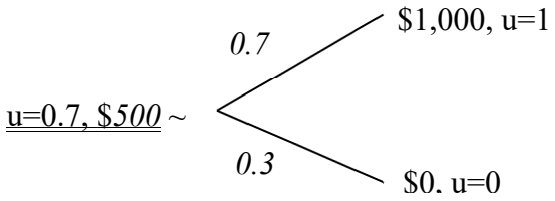
This means that $0 \le u(x_i) \le 1$ for all outcomes involved in the decision.

To find $u(x)$ for other values x in the interval from *Lo* to *Hi*, we can follow two alternative methods of asking questions, both based upon indifference statements involving simple lotteries.

1. **Value approach**: Given $p = P[A]$ and any two of the amounts *A*, *B*, *C*, and their corresponding utilities, find the third amount and its utility. This is in fact the best method: select *A* and *C* and an easy-to-understand *p* (such as *p*=50%) and ask the decision maker for *B*.

2. **Probability approach**: Given the three amounts *A*, *B*, *C*, and two of their utilities, find *P*[*A*] and the utility of the third amount.

Example 1. Helen is faced with a lottery that pays *A*=\$1000 with probability 0.7 and *C*=\$0 otherwise. She states that she would sell it for no less than *B*=\$500: i.e., \$500 ~ (\$1,000, 0.7; \$0).

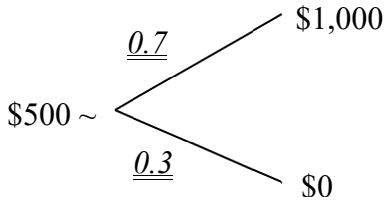


Suppose that we assign u(\$1,000)=1 and u(\$0)=0:

$$u(\$500) = u(\$1{,}000)\,(0.7) + u(\$0)\,(0.3)$$

$$u(\$500) = (1)\,(0.7) + (0)\,(0.3) = 0.7$$

Example 2. Helen is asked to give up the lottery (\$1,000, *p*; \$0) for a given sure payoff of \$500. Helen responds that she would not do so if *p* >= 0.7. This statement implies the same lottery as above and also leads to u(\$500) = (1) (0.7) + (0) (0.3) = 0.7



Note that the probability approach is not an easy method to use in practice because most people cannot distinguish easily between different values of probability. Thus, this method requires the use of visual aids, such as a *probability wheel*, to help the decision maker visualize probabilities, such as 70%.

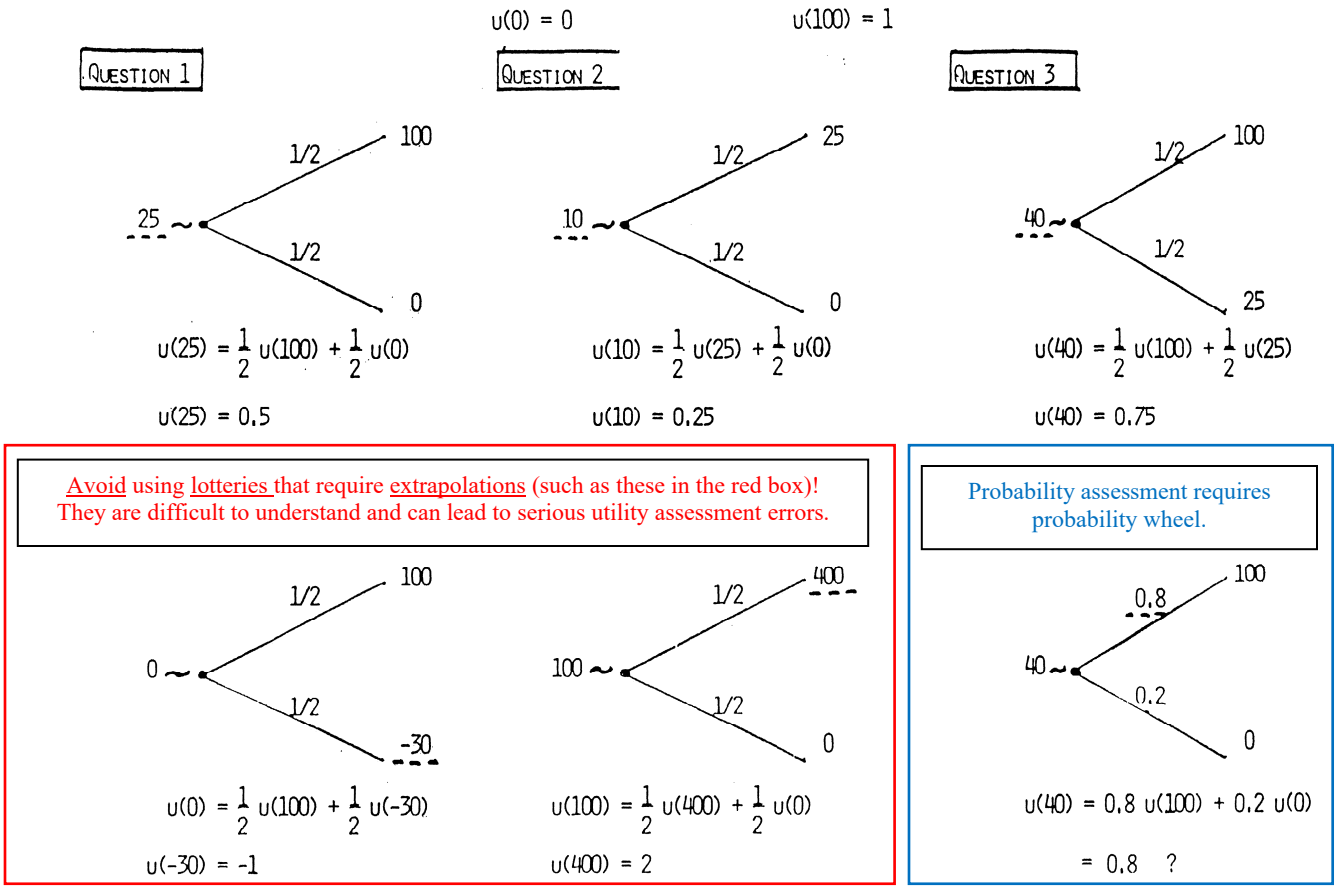
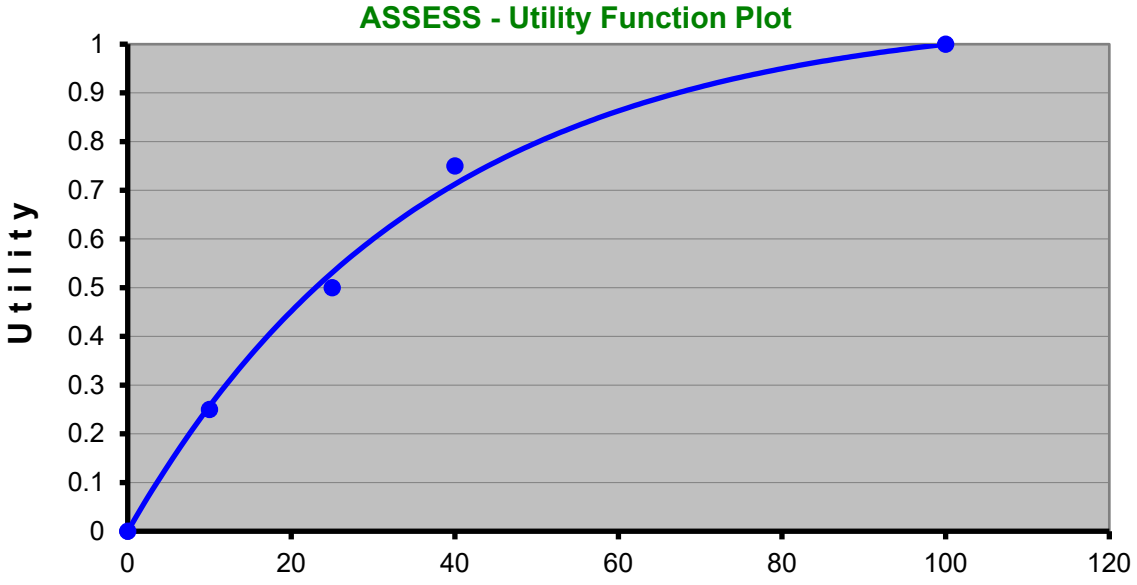


FIGURE 20 RESPONSES TO UTILITY ASSESSMENT QUESTIONS.



Caution: Utilities cannot be added

Utility is like temperature. We cannot add the utilities of two outcomes D and E to get the utility of the sum of the two outcomes, D+E:

$$u(D)+u(E) \neq u(D+E)$$

For example: $u(500)+u(500)=0.7+0.7=1.4 \quad \neq \quad u(500+500)=u(1{,}000)=1$.

Caution: Utility Assessment should NOT use lotteries with relatively small gambles

One frequent "error" in judgment is that people are very unreliable in their judgment about small gambles (Schlaifer 1962, Bell 2007).

A person might say

I would exchange a 50–50 gamble between \$0 *and* \$10 *for* \$4.

The same person might say

I would exchange a 50–50 gamble between \$0 *and* \$100,000 *for* \$35,000.

No plausible utility function can capture a person's risk attitude
both "in the small" (tens of dollars) and "in the large" (tens of thousands of dollars).

Based on his practical experience, Schlaifer warns that it is inadvisable to rely on any utility assessments that were not in terms of sums of money of the order of the respondent's assets.

So, if a person's total assets are in the hundreds of thousands of dollars, then the assessment of that person's utility function should use lotteries whose prizes also differ by hundreds of thousands of dollars.

Bell, D.E. (2007). "Utility and Risk Preferences", in *Advances in Decision Analysis From Foundations to Applications*, Ed: Edwards W., Miles R.F., von Winterfeldt, D. Cambridge University Press, Cambridge CB2 8RU, UK

Schlaifer, R. O. (1962). "Assessment of the utility of money". Unpublished seminar notes. Harvard University, Graduate School of Business Administration.

Risk Averse or Gambler?

A person is considered risk averse if he would avoid an even (50-50) chance of doubling his money. Conversely, a risk-preferring person (a gambler) would accept less than an even chance to double his money. A person may be averse to risk or desire risk at two different times or with regard to two different ranges in money.

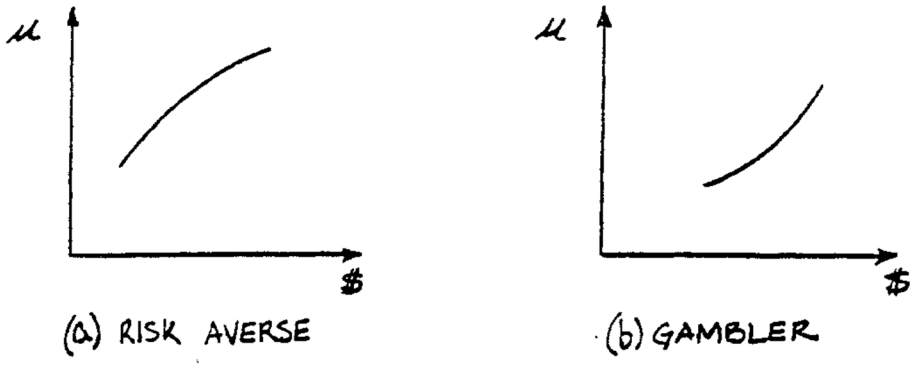


Figure U.1. Risk Averse and Gambler Utility Curves

The degree to which a person is risk-averse or risk-preferring (a gambler) is directly related to the *curvature* of his/her utility function over the range of possible risky outcomes. When the range of outcomes is small compared to the overall range of possible money decisions for which the utility function is defined, the utility function may be assumed to be linear. This is similar to saying that over a small interval dx any function can be approximated well by the tangent or the secant to the function in that interval.

Mixed-behavior Utility Functions

Some people argue that the figure below shows a possible utility function for a construction company where the decision maker has limited personal liability.

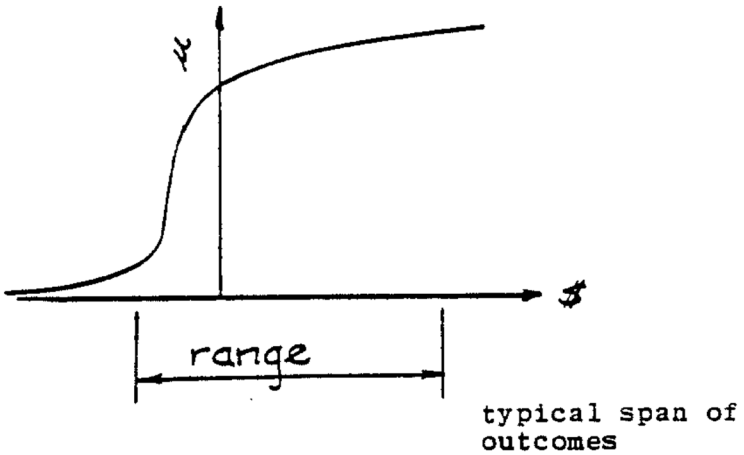


Figure U.2. Utility Curve Exhibiting Both Risk-Averse and Gambler Behavior

This utility function implies that for decisions that involve positive profit outcomes the contractor is risk averse. For outcomes that represent significant losses, however, the behavior changes to that of a gambler.

This is an illustration of a general phenomenon discovered by psychologists Kahneman and Tversky where people tend to be risk-averse when dealing with profit and risk-seeking when deciding about loss. One could argue that when a company is faced with risks that may drive it into bankruptcy, it may wish to take chances that it would otherwise avoid so that it may simply have a chance to survive. Although possible, this behavior need not apply to all contractors.

The overall shape of a utility function may not vary significantly from one risk-averse person to the next. What varies a lot from one person to the next is the range of monetary outcomes (dollars) that is necessary for two utility functions to exhibit a similar degree of curvature (risk-aversion). People whose net worth is small, who are very risk averse, and who deal with decisions involving small amounts of money have utility functions that show a lot of curvature over a narrow range of dollar outcomes. In contrast, people with a large net worth, who are less risk averse, and who deal with decisions that involve very large amounts of money, need a much wider range of dollar outcomes to exhibit the same degree of curvature in their utility functions.

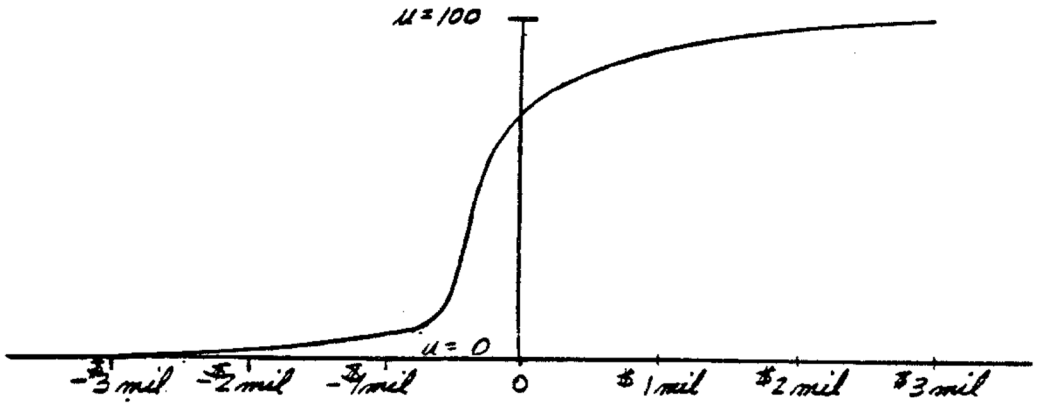


Figure U.3. Effect of Monetary Scale on Utility Curve

For example, the above utility function might be a typical curve for a contractor with about $200,000 in working capital. Dividing or multiplying all numbers along the *x*-axis by 10 would give typical curves for working capital = $20,000 or $2,000,000 respectively.

It is typical within a company for employees at different levels of management to be using similar curves in their decisions but with the ranges dependent upon their levels within management. A utility function may be defined for each separate decision or for all decisions depending on its range of outcomes. A utility function may be used for any decision problem whose outcomes are within the function's range. The reverse is not true: A utility function cannot be used for a decision whose range of outcomes is outside the range of the utility function.

It is important that all decisions within a company use the same curve. Such a curve can be defined by top management and made available for important decisions. Because any curve is essentially linear over ranges that are small compared to its overall range, lower level decisions should probably use a linear function, i.e., actual dollars.

Expected Utility Value

The methodology that we use for calculating expected monetary value (EMV) is equally effective in analyzing decisions based upon nonlinear utility functions. We will calculate the expected utility value (EUV) just as we have calculated EMV and we will prefer whichever alternative has the highest EUV. The use of EUV is somewhat more complex than EMV because (1) a utility curve must be developed and (2) its nonlinearity requires care in calculations. This is a small price to pay, however, for the increased accuracy and applicability of EUV for important decisions.

Use of EUV in selecting among alternatives follows these steps:

1. List the alternatives which are to be evaluated.
2. List the sample set of possible chance outcomes for each alternative, estimate the total monetary value (or other unit of value, such as grade) of each outcome, and estimate the probability of each outcome.
3. Develop a utility function u(x) which will convert the total estimated value of each outcome to a measure of its utility.
4. Find the utility $u(x_i)$ for each outcome x_i.
5. Calculate the EUV of each alternative using the same equation as for EMV:

$$EUV = E[u(x)] = \sum_i u(x_i)P[x_i]$$

6. Select the alternative with the highest EUV.

Example. The example problem described earlier in Fig. E.7 can be used to demonstrate EUV calculations. The range in monetary value of outcomes is from \$462,500 to -\$322,700. Fig. U.4 shows example utility curves which will be used to evaluate the decision alternatives, one curve for a risk-averse decision maker and one curve for a gambler.

The decision evaluation for the risk-averse curve is shown in Fig. U.5. The utility of each outcome is found from the risk-averse curve and is used to calculate EUV at each chance node. The branch with the largest EUV is selected at each decision node. The small plant is preferred over the larger plant and, like the decision for EMV, it will not be expanded if the demand in the first two years is low, and will be expanded if the demand in the first two years is high.

The decision for the gambler's utility curve is shown in Fig. U.6. We can expect that this curve represents a company in trouble, willing to take a risk to get out of a bad situation. Its utilities are different from that of the risk-averse person and so is its decision preference. Whereas the straight-line person represented by EMV and the risk-averse person would select the small plant and later expand if indicated, the gambler starts off big, purchasing the large plant.

Also, notice that the amounts of money equivalent to the EUV are different for the two utility curves, and both are different from the EMV = \$289,200. The EUV of 89.8 for the risk averse curve is equivalent to about \$245,000 and the EUV of 66.2 for the gambler is equivalent to about \$333,000 on his curve. Thus, a risk-neutral person choosing by EMV would be indifferent between a choice of \$289,200 for sure and the investment opportunity. The risk-averse person would be indifferent at \$245,000 and would give up the opportunity if presented with an amount above that figure. The gambler would demand more than \$333,000 to give up his opportunity.

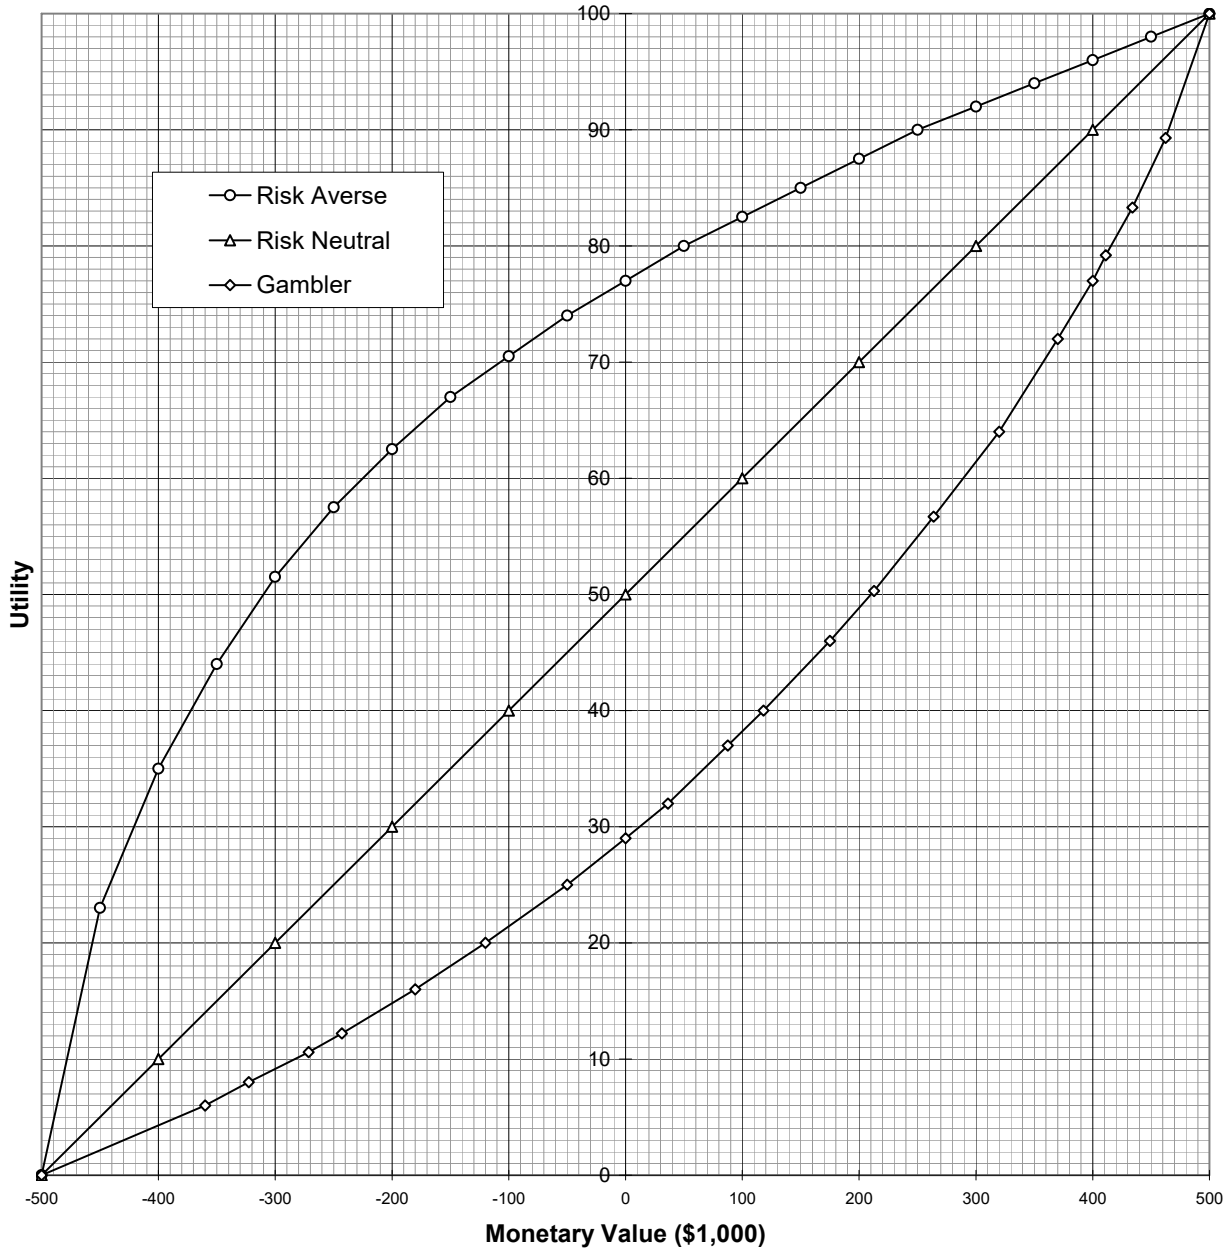


Figure U.4. Utility Curves for Plant Decision Problem

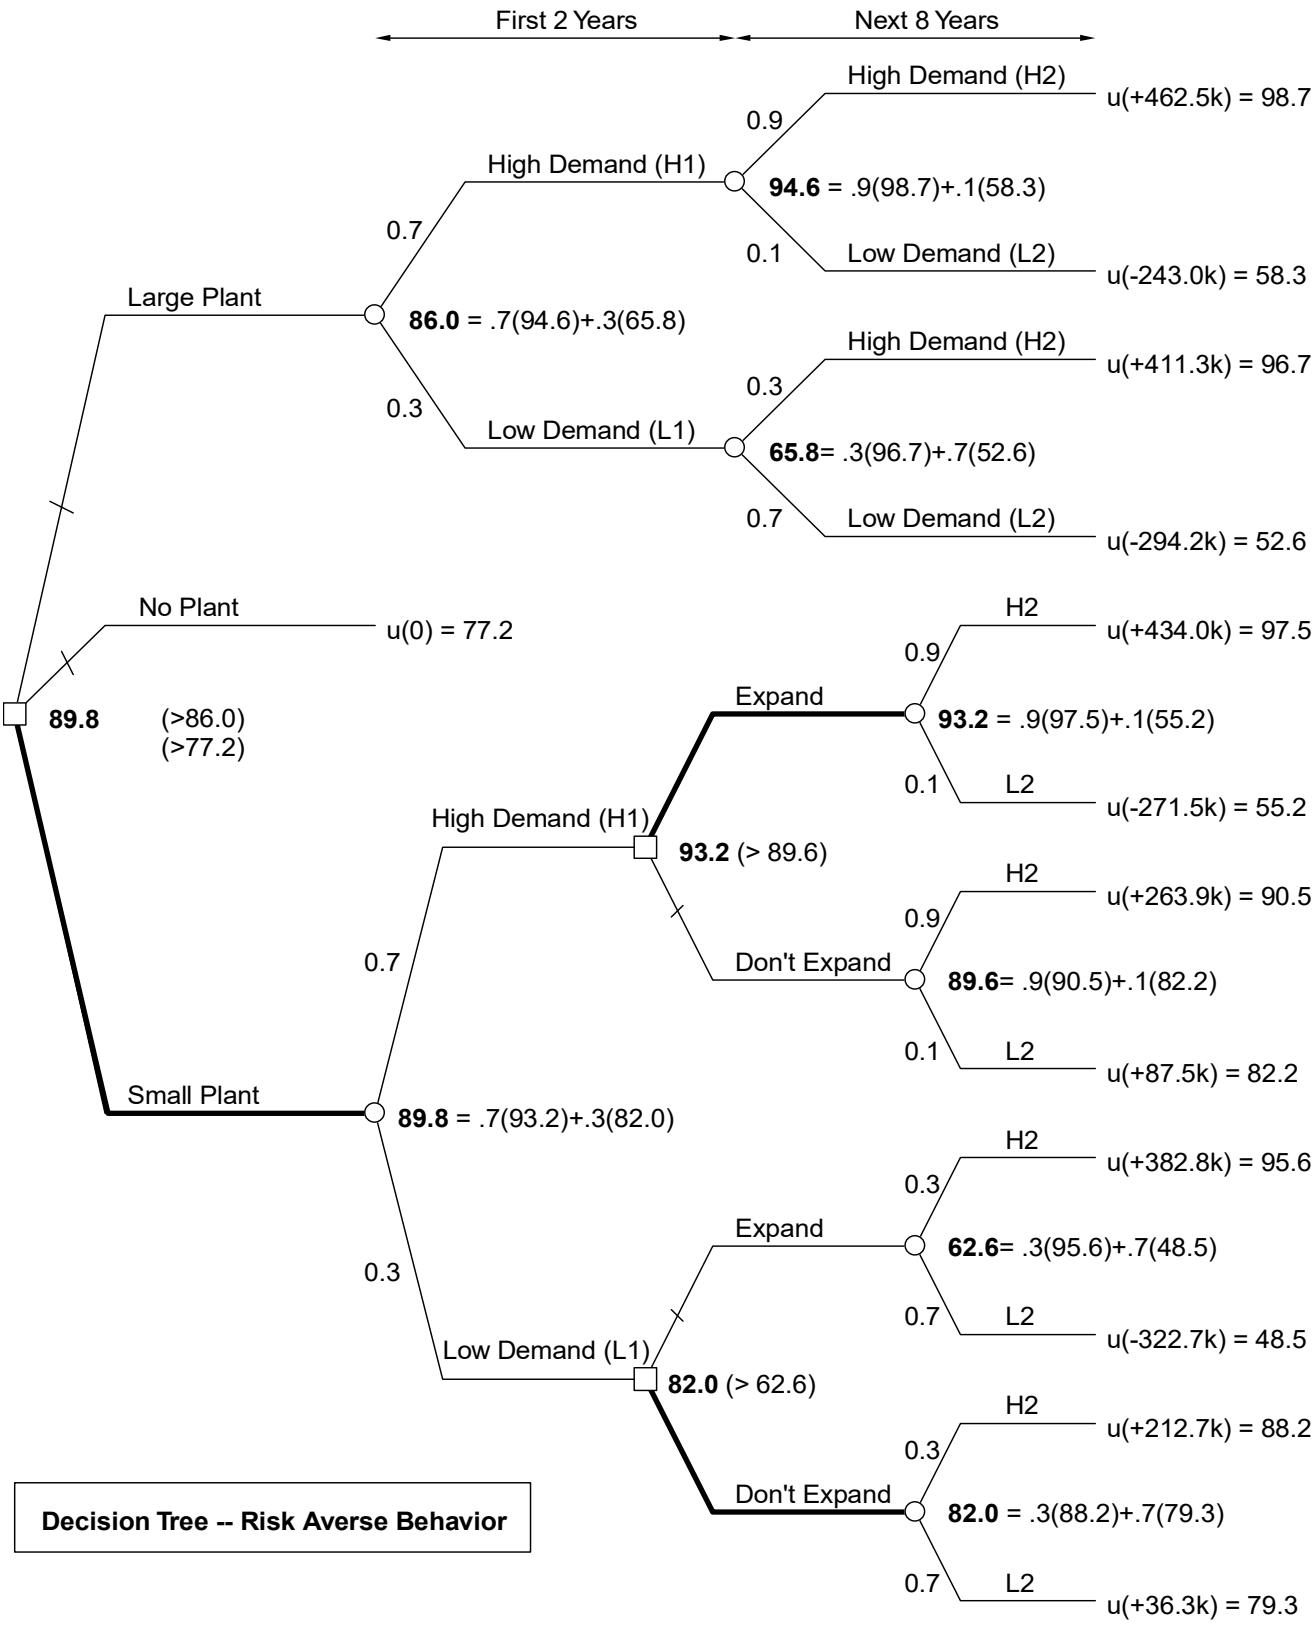


Figure U.5. Expected Utility Solution for Plant Decision Problem—Risk Averse Behavior

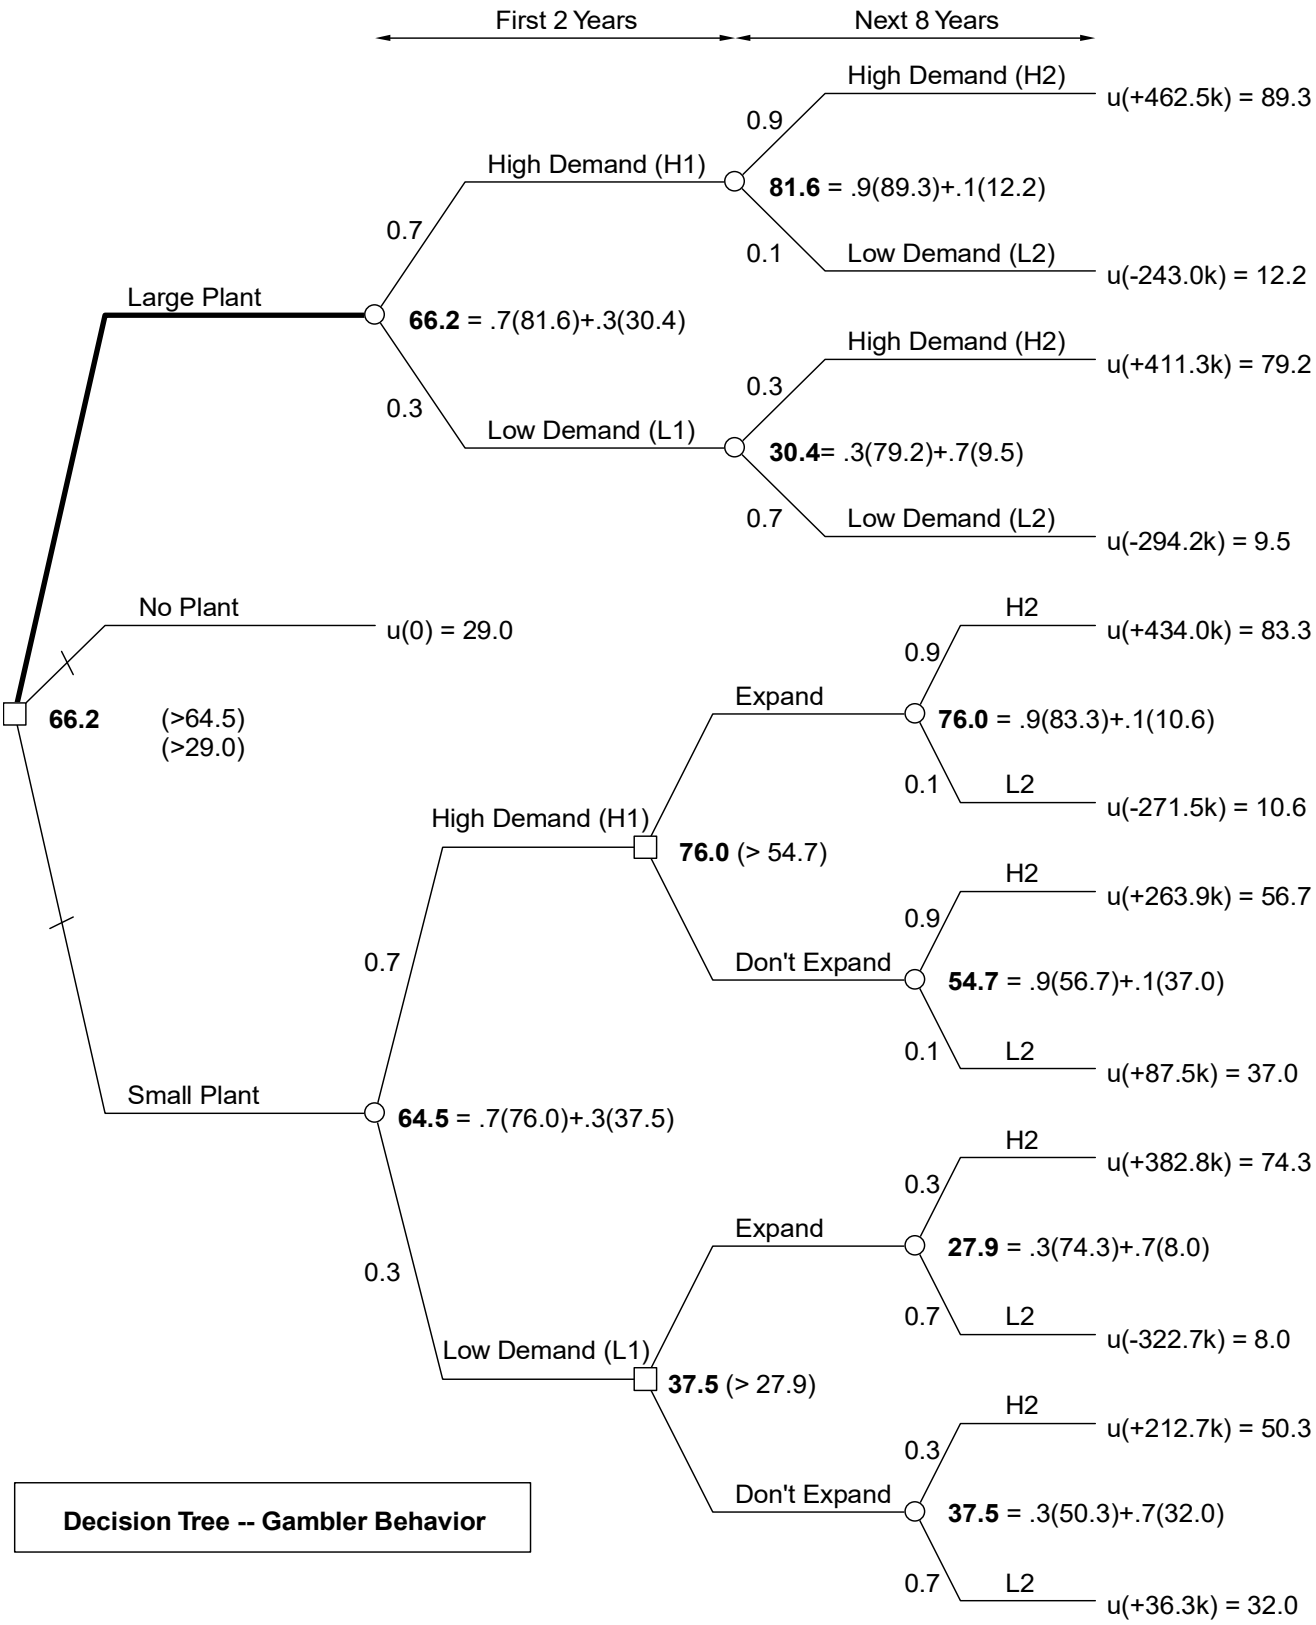


Figure U.6. Expected Utility Solution for Plant Decision Problem—Gambler Behavior

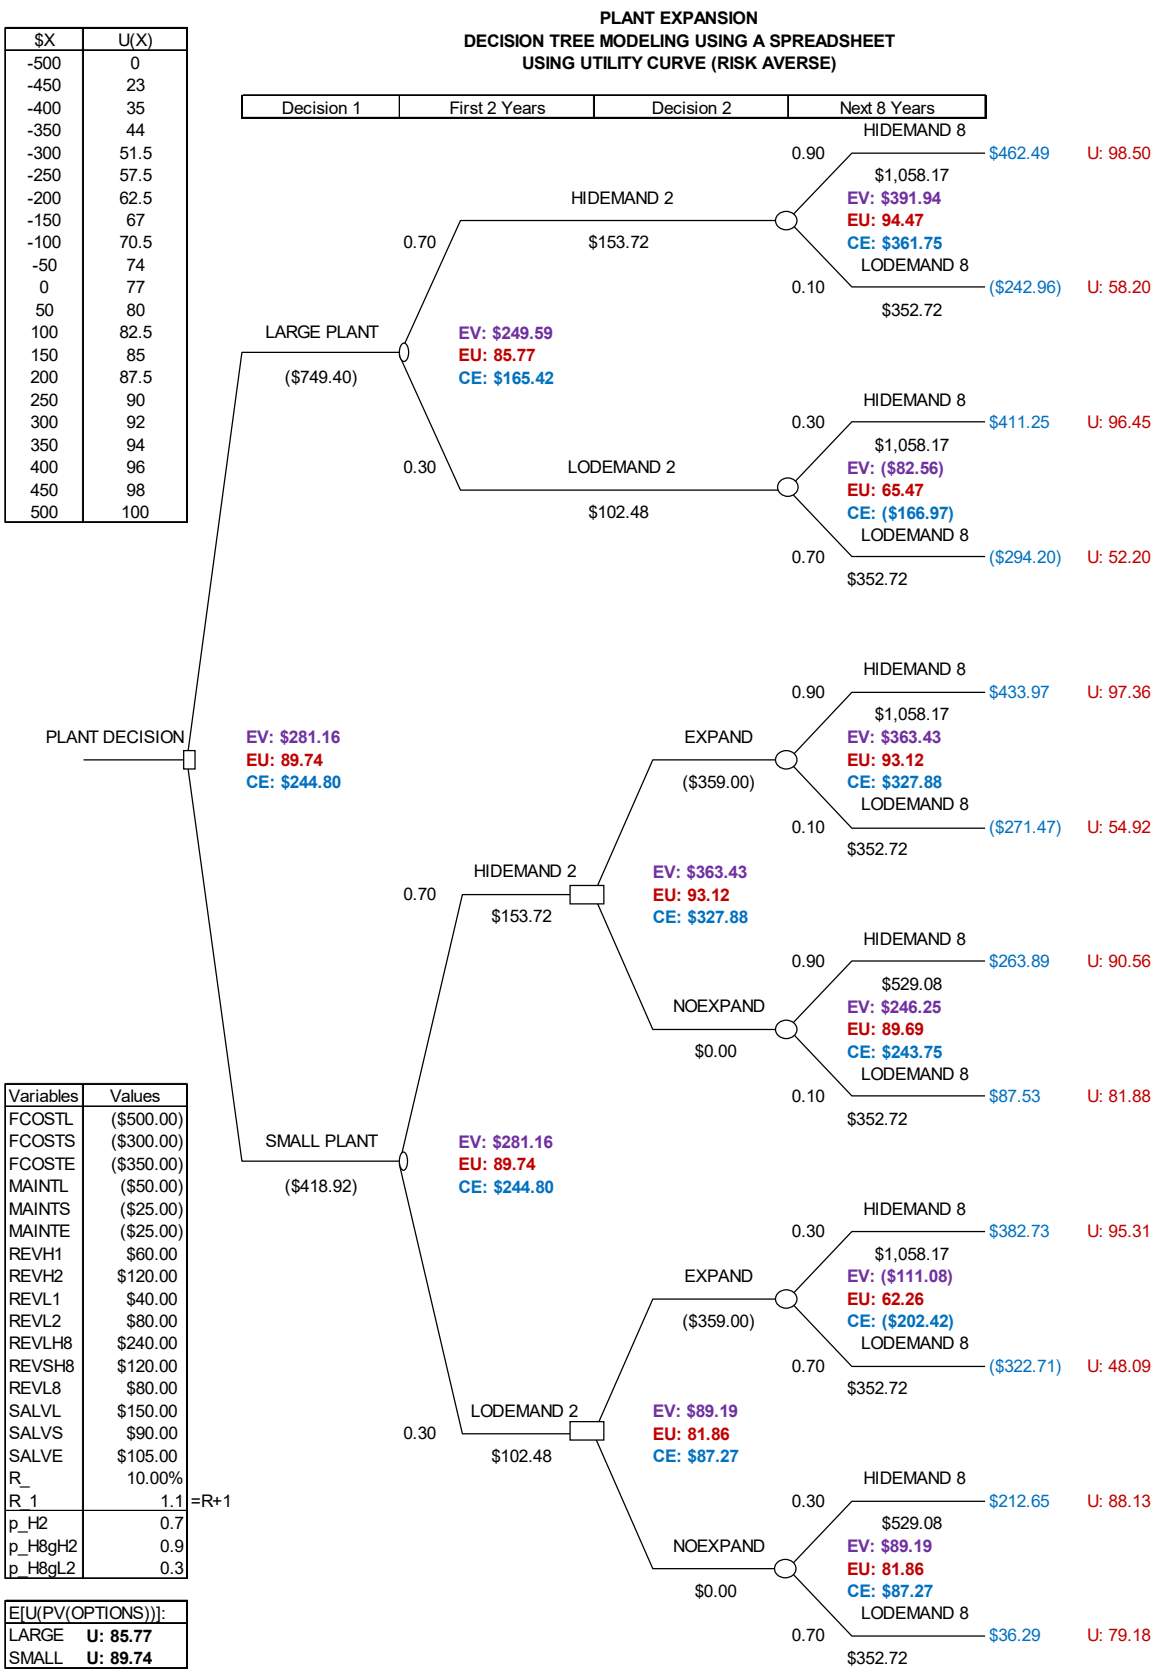


Figure U.7. Parametric Expected Utility Solution for Plant Example—Risk-Averse Behavior

PLANT EXPANSION
DECISION TREE MODELING USING A SPREADSHEET
SENSITIVITY ANALYSIS
USING UTILITY CURVE (RISK AVERSE)

One-way Sensitivity Analysis

p_H2	Large Plant E[U]	Small Plant E[U]	Best
0	65.47	81.86	S
0.1	68.37	82.99	S
0.2	71.27	84.11	S
0.3	74.17	85.24	S
0.4	77.07	86.36	S
0.5	79.97	87.49	S
0.6	82.87	88.61	S
0.7	85.77	89.74	S
0.8	88.67	90.86	S
0.9	91.57	91.99	S
1	94.47	93.12	L

Two-way Sensitivity Analysis

p_H2	p_H8gL2										
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	S	S	S	S	S	S	S	S	L	L	L
0.1	S	S	S	S	S	S	S	S	L	L	L
0.2	S	S	S	S	S	S	S	S	L	L	L
0.3	S	S	S	S	S	S	S	S	L	L	L
0.4	S	S	S	S	S	S	S	S	L	L	L
0.5	S	S	S	S	S	S	S	S	L	L	L
0.6	S	S	S	S	S	S	S	S	L	L	L
0.7	S	S	S	S	S	S	S	L	L	L	L
0.8	S	S	S	S	S	S	S	L	L	L	L
0.9	S	S	S	S	S	L	L	L	L	L	L
1	L	L	L	L	L	L	L	L	L	L	L

p_H8gH2	Large Plant E[U]	Small Plant E[U]	Best
0	60.38	81.87	S
0.1	63.21	82.48	S
0.2	66.03	83.09	S
0.3	68.85	83.70	S
0.4	71.67	84.30	S
0.5	74.49	84.91	S
0.6	77.31	85.52	S
0.7	80.13	86.13	S
0.8	82.95	86.77	S
0.9	85.77	89.74	S
1	88.59	92.71	S

p_H8gH2	p_H2										
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	S	S	S	S	S	S	S	S	S	S	S
0.1	S	S	S	S	S	S	S	S	S	S	S
0.2	S	S	S	S	S	S	S	S	S	S	S
0.3	S	S	S	S	S	S	S	S	S	S	S
0.4	S	S	S	S	S	S	S	S	S	S	S
0.5	S	S	S	S	S	S	S	S	S	S	S
0.6	S	S	S	S	S	S	S	S	S	S	S
0.7	S	S	S	S	S	S	S	S	S	S	S
0.8	S	S	S	S	S	S	S	S	S	S	L
0.9	S	S	S	S	S	S	S	S	S	S	L
1	S	S	S	S	S	S	S	S	S	S	L

p_H8gL2	Large Plant E[U]	Small Plant E[U]	Best
0	81.79	88.93	S
0.1	83.12	89.20	S
0.2	84.44	89.47	S
0.3	85.77	89.74	S
0.4	87.10	90.01	S
0.5	88.43	90.28	S
0.6	89.75	90.55	S
0.7	91.08	90.81	L
0.8	92.41	91.08	L
0.9	93.74	92.36	L
1	95.06	93.77	L

p_H8gL2	p_H2										
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	S	S	S	S	S	S	S	S	S	S	L
0.1	S	S	S	S	S	S	S	S	S	S	L
0.2	S	S	S	S	S	S	S	S	S	S	L
0.3	S	S	S	S	S	S	S	S	S	S	L
0.4	S	S	S	S	S	S	S	S	S	S	L
0.5	S	S	S	S	S	S	S	S	S	L	L
0.6	S	S	S	S	S	S	S	S	S	L	L
0.7	S	S	S	S	S	S	S	L	L	L	L
0.8	L	L	L	L	L	L	L	L	L	L	L
0.9	L	L	L	L	L	L	L	L	L	L	L
1	L	L	L	L	L	L	L	L	L	L	L

Figure U.8. Sensitivity Analysis for Plant Example—Risk-Averse Behavior

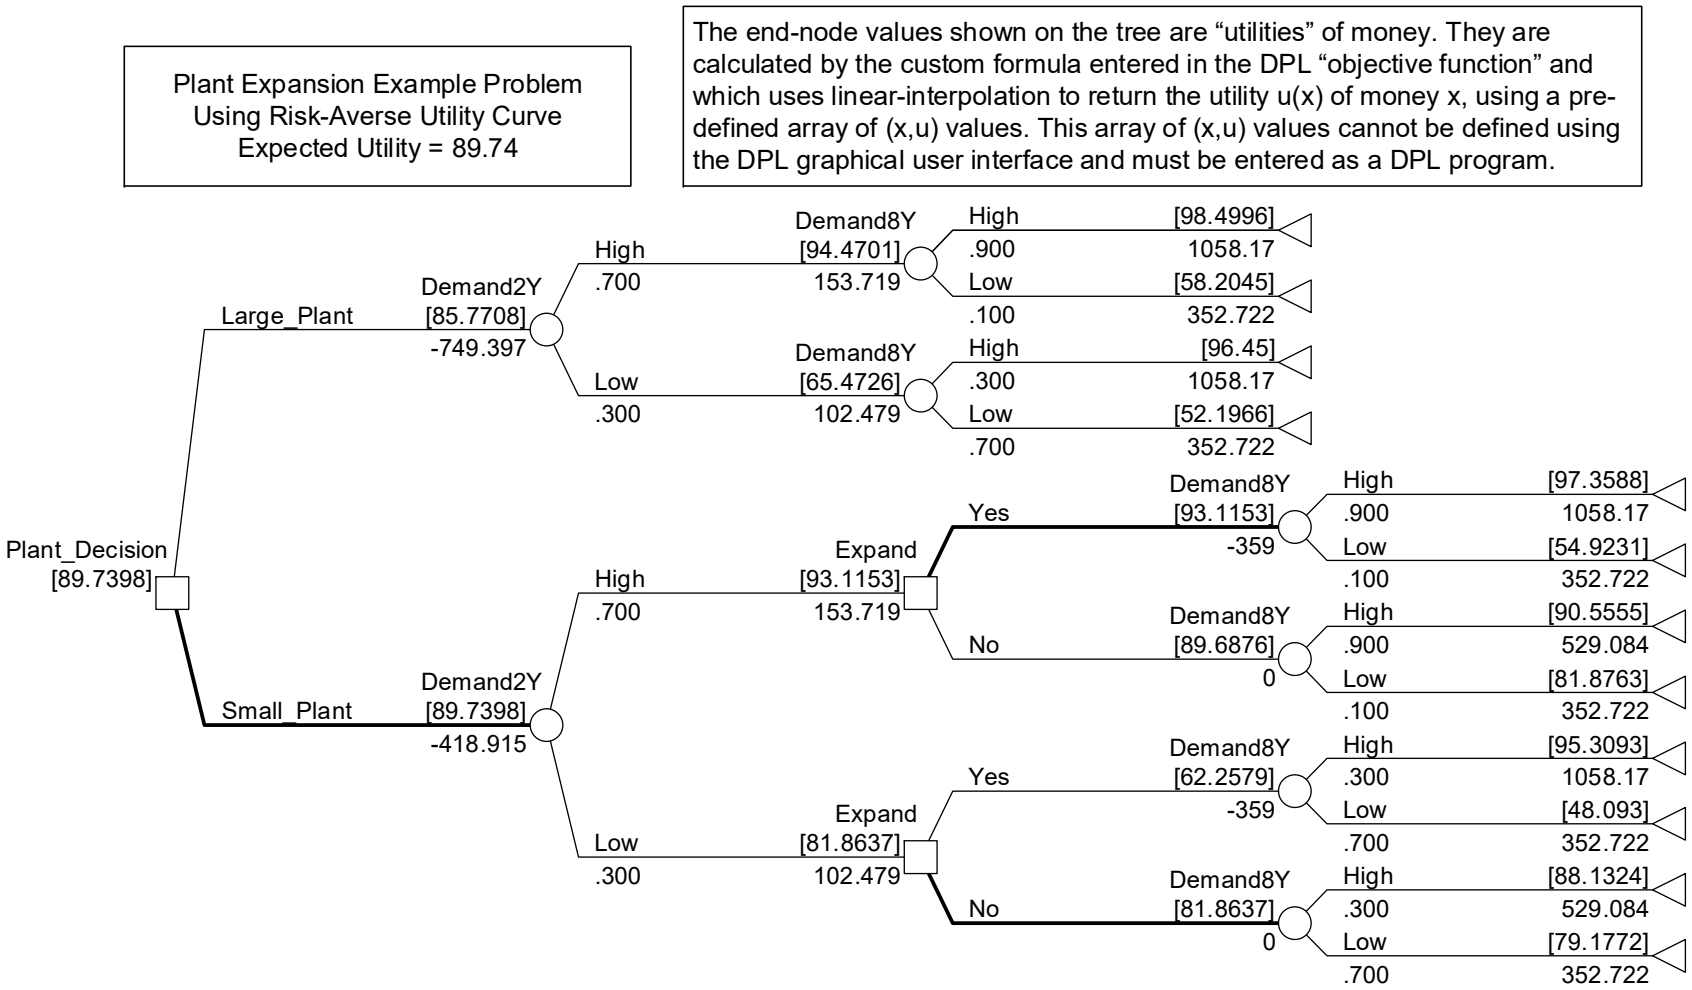


Figure U.9. Expected Utility Solution in DPL for Plant Example—Risk-Averse Behavior

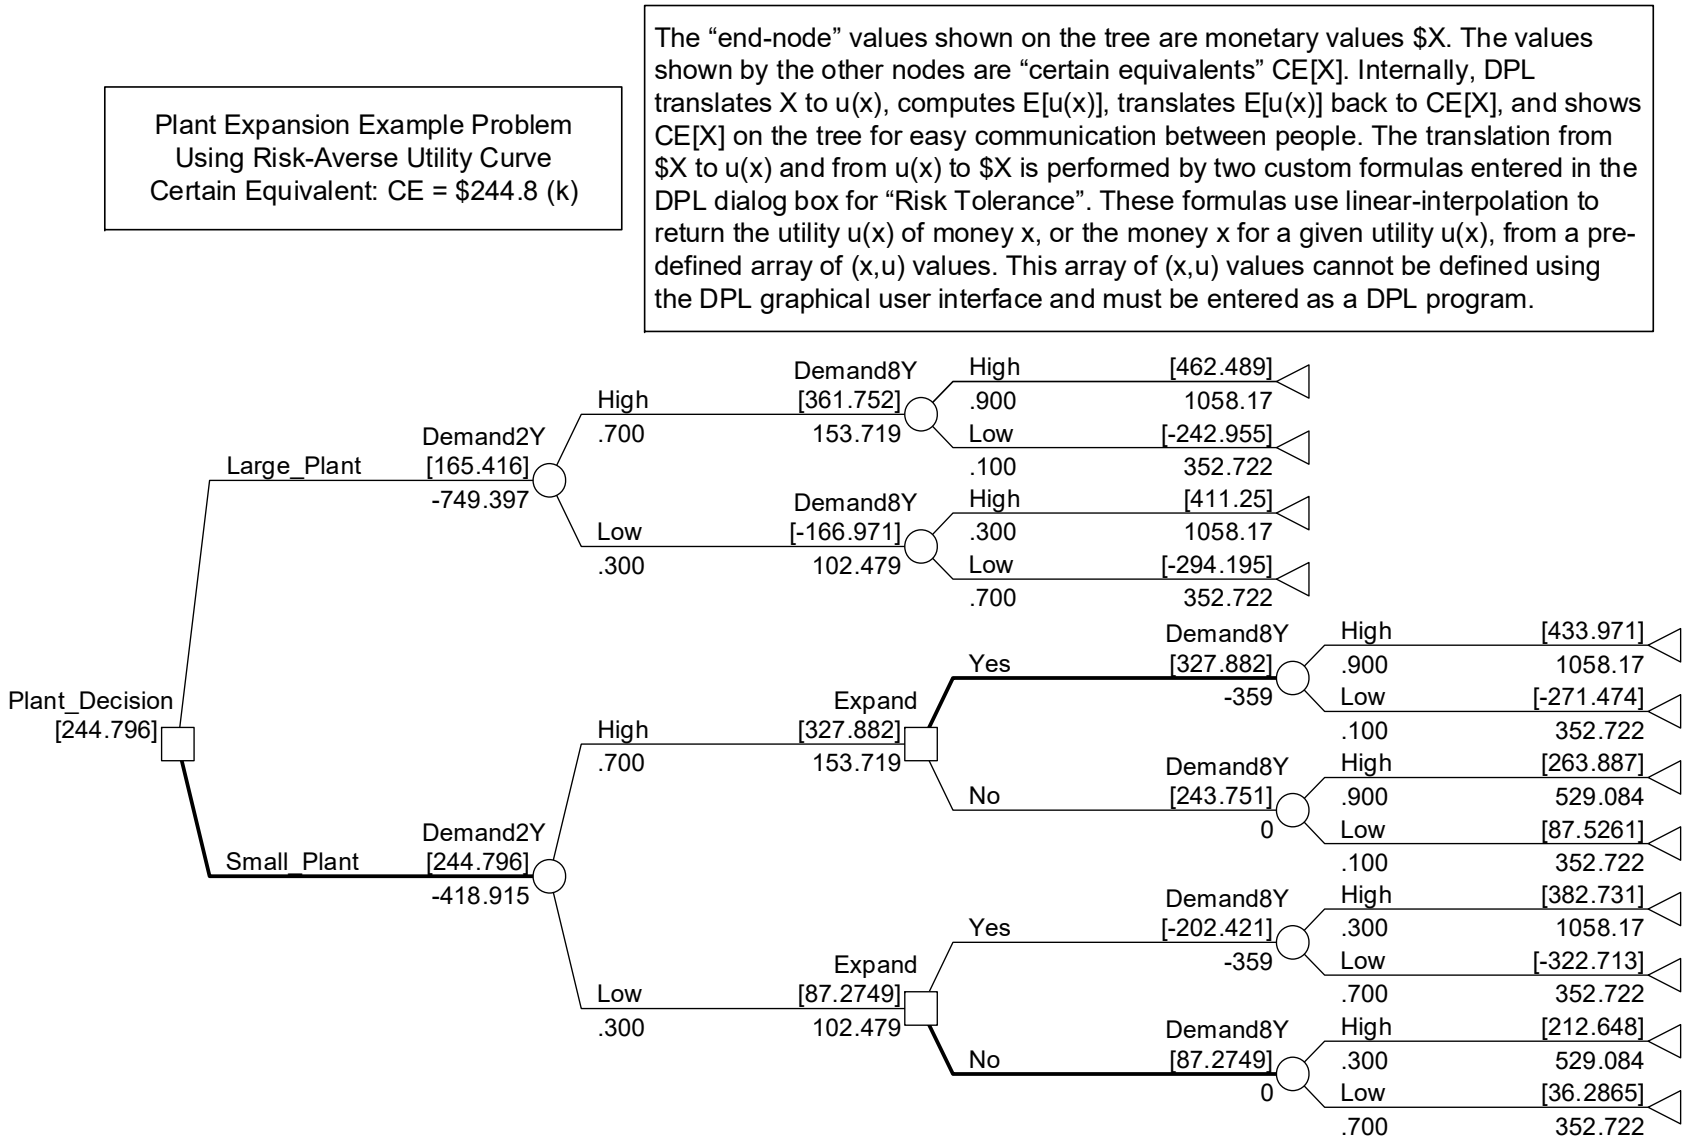
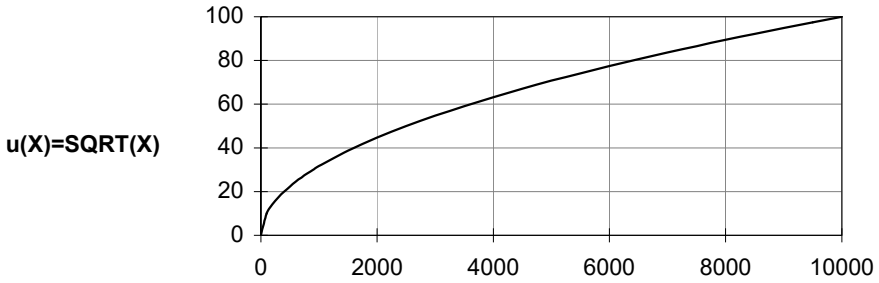


Figure U.10. Certain Equivalent Solution in DPL for Plant Example—Risk-Averse Behavior

THE NON-LINEARITY OF VALUE AND THE IMPORTANCE OF THE NET ASSET POSITION.



Case	E[u]	CE[X]	$X	u(X)
A	30	900 ~	1/2 400; 1/2 1600	20; 40

Case	E[u]	CE[X]	$X	u(X)
B	75	5625 ~	1/2 3600; 1/2 8100	60; 90

A+B	81.5	6644 ~	1/4 400 + 3600; 1/4 1600 + 3600; 1/4 400 + 8100; 1/4 1600 + 8100	63.2; 72.1; 92.2; 98.5

The sum of CE[A] and CE[B] does not equal CE[A+B]

900
5625
6525 <> 6643.9

B+CE[A]	81.5	6644 ~	1/2 3600 + 984; 1/2 8100 + 984	67.7; 95.3

Selling price of A, when owning A & B and keeping B, goes up (requires trial & error):

984.4 > 900

Notice: outcomes of B increase by Delta = 984.4 but CE[B] goes up by 6644-5625 = 1019 > 984.4

B|A=400 CE[B]=?	77.7	6040 ~	1/2 3600 + 400; 1/2 8100 + 400	63.2; 92.2

Selling price for B given A=$400

6040
-400
CE[B|A=900]= 5640 > 5625

B|A=1600 CE[B]=?	85.3	7276 ~	1/2 3600 + 1600; 1/2 8100 + 1600	72.1; 98.5

Selling price for B given A=$1600

7276
-1600
CE[B|A=1600] 5676 > 5640

CE[B]+A =A+B CE[B]=?	81.5	6644 ~	1/2 400 + 5657; 1/2 1600 + 5657	77.8; 85.2

Selling price for B when owning both A & B and keeping A (by trial and error)

> 5640
CE[B] = 5657
< 5676

When faced with sequential lotteries (e.g., A+B) we cannot a substitute any single lottery with its certain equivalent (e.g., A+B <> A+CE[B])

The CE[B] depends on our current asset position. It is lowest when we own only B. If we own both A+B then CE[B] depends on what we know about A's outcome (5625 | no A; 5640 | A=400; 5657 | A; 5676 | A=1600)

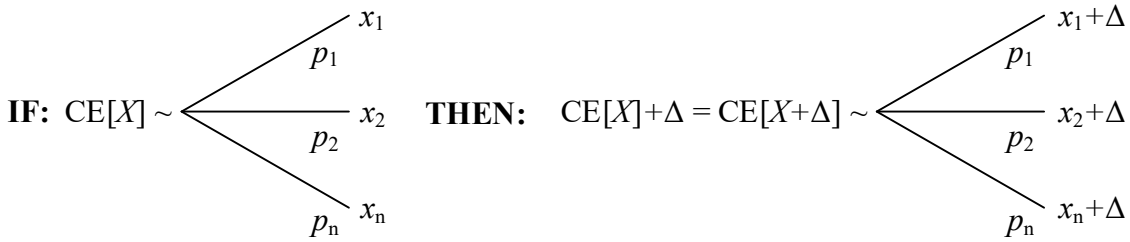
Delta Property—Constant Risk Preference—Exponential Utility

The "*delta-property*" can be defined as follows:

- "If all *prizes* of a lottery are increased by the same amount Δ,
  then the *certain equivalent* of the lottery is also increased by the same amount Δ."

$$CE[X+\Delta] = CE[X] + \Delta$$

- The *delta-property* is illustrated graphically by the following *indifference* statements:



Acceptance of the Delta Property restricts the utility curve to be either linear or exponential:

$$\text{Linear: } u(x) = a + bx$$

$$\text{Exponential: } u(x) = a - be^{-\gamma x}$$

Notice that the exponential utility function is completely defined by just one parameter, the *risk aversion coefficient* γ, or its reciprocal, the *risk tolerance* $\rho = 1/\gamma$:

	γ = risk aversion coefficient	$\rho = 1/\gamma$ = risk tolerance	
Risk Averse Behavior:	$\gamma > 0$	$\rho > 0$	$b > 0$
Risk Preferring (*Gambler*):	$\gamma < 0$	$\rho < 0$	$b < 0$
Risk Neutral:	$\gamma \to 0$	$\rho \to \infty$	

A person who accepts the *delta property* is often referred to as a "*delta person*". He is also called a "*constant risk averter*" or a "*person with constant risk aversion*" because the *risk premium* for a lottery does not change if all lottery prizes are increased by the same amount Δ.

The exponential utility curve (like linear utility) has some very useful properties:

- The selling price and the buying price of a lottery are identical.
- When a delta person is engaged in a sequence of lotteries, his certain equivalent for the whole prospect is the sum of the certain equivalents for each individual lottery.
- For a lottery X described by a Normal distribution with mean m and variance σ^2 then the certain equivalent of the lottery is given in closed form by the very useful formula:

$$CE[X] = m - \frac{1}{2}\gamma\sigma^2$$

The usefulness of the *delta property* in practice cannot be overemphasized. It makes calculations so *significantly easier*, that often (as a first approximation) it is assumed that the decision maker is indeed a delta-person to allow a much easier preliminary investigation and sensitivity analysis.

Value of Information = PIBP for that Information

Let CT = the *Cost of the Test* by which to obtain information.

The *Value of Information* (perfect or sampled) is that *special* cost, CT^*, that makes the decision maker *indifferent* between obtaining and not obtaining the information. Mathematically:

$$E[u(X\text{-}CT^*) \mid \text{info}, CT^*] = \max E[u(X) \mid \text{no info}]$$

Or, equivalently,

$$CE[X\text{-}CT^* \mid \text{info}, CT^*] = \max CE[X \mid \text{no info}]$$

Thus, the *Expected Value of Perfect Information* (*EVPI*) equals the *Personal Indifferent Buying Price* (*PIBP*) for that perfect information.

$$EVPI = CT_{PI}^* = PIBP \text{ for that perfect information}$$

And the *Expected Value of Sampled Information* (*EVSI*) equals the *Personal Indifferent Buying Price* (*PIBP*) for that sampled information.

$$EVSI = CT_{SI}^* = PIBP \text{ for that sampled information}$$

Of course, because EVPI>=EVSI

$$CT_{PI}^* >= CT_{SI}^*$$

Finding the *EVPI* or the *EVSI* when using a nonlinear utility curve and thus selecting alternatives based on *Expected Utility Value* (*EUV*) requires trial-and-error and can be time-consuming.

The trial-and-error procedure starts by assigning an initial estimate, CT, of the cost in dollars for the decision tree branch (let us call it the *test* branch) associated with obtaining the information.

This trial cost CT is then subtracted from the dollar values at all the end-nodes of the test subtree.

The utility of each outcome is then found, and the *EUV* of the test subtree is calculated and compared to the highest *EUV* without the test. If the *EUV* of the test subtree at the trial cost CT is the same as the highest *EUV* without the test, then $CT = CT^*$.

Otherwise, we try a lower or higher value for CT. Our objective is to find two values of CT that (a) bracket CT^*, and (b) are close enough to each other, so that we can interpolate CT^* reliably.

Special Cases:

When selecting alternatives based on *Expected Monetary Value* (*EMV*), we have:

$$CE[X\text{-}CT^* \mid \text{info}, CT^*] = E[X\text{-}CT^* \mid \text{info}, CT^*] = E[X \mid \text{info}, CT=0] - CT^* = \max CE[X \mid \text{no info}]$$

$$CT^* = E[X \mid \text{info}, CT=0] - \max E[X \mid \text{no info}]$$

Thus, in the special case of a linear utility curve (in which case the decision maker is Risk-Neutral and makes decisions based on *EMV*), we can find CT^* as the difference between the *EMV with the information* minus the *best EMV without the information*.

As illustrated later, a similar result applies for the exponential utility curve because of the *delta property*. The only difference is that $E[\ldots]$ is changed to the more general $CE[\ldots]$.

$$CT^* = CE[X \mid \text{info}, CT=0] - \max CE[X \mid \text{no info}]$$

Example—Expected Value of Perfect Information (EVPI).

We shall now calculate the *Expected Value of Perfect Information* (*EVPI*) for the prefabricated plant decision example when using expected utility values (EUV) and the risk-averse utility curve shown in Fig. U.4.

For reference, the decision tree for calculating EVPI when using EMV is shown in Fig. E.18 and gave a value of:

$$\text{EVPI}=\$60{,}900.$$

Fig. U.5 and Fig. U.7 show the decision tree without perfect information when using EUV and the risk-averse utility curve. From these trees we have:

$$\max E[u(X) \mid \text{no info}] = 89.8$$

$$\max CE[X \mid \text{no info}] = \$244.8\text{k}$$

Fig. U.11 and Fig. U.12 show the *test subtree* when using EUV to calculate the EVPI for perfect information about the "Demand Over Ten Years".

Each end-node of the test subtree in Fig. U.11 and Fig. U.12 is a rearranged end-node value from Fig. U.5 and Fig. U.7 minus the cost of the research, CR or CT_10, that would provide perfect information.

Fig. U.11 illustrates the trial-error procedure. A trial cost of CR=\$100,000 was used first, producing EUV = 88.8. This is less than the EUV of 89.8 shown on Fig. U.5 without the test. A second trial value of CR=\$80,000 produced EUV = 89.6 and a third trial value of CR=\$70,000 produced EUV = 90.15. Thus, $\$70\text{k} \le CR^* \le \$70\text{k}$. Linear interpolation produced:

$$\text{EVPI}_{\text{EUV}}=CR^*=\$76{,}000 \;>\; \text{EVSI}_{\text{EMV}} = \$60{,}900$$

Fig. U.12 shows how to automate the trial-and-error procedure by using the "goal seek" capability of Excel which tries to make the cell DeltaU equal to 0 by changing the cell CT_10. The calculated EVPI is slightly different due to the accuracy of the two methods:

$$\text{EVPI}_{\text{EUV}}=CT_10^*=\$74{,}128 \;>\; \text{EVSI}_{\text{EMV}} = \$60{,}900$$

This illustrates that a risk-averse decision maker is willing to pay more for perfect information that reduces uncertainty than a risk-neutral decision maker.

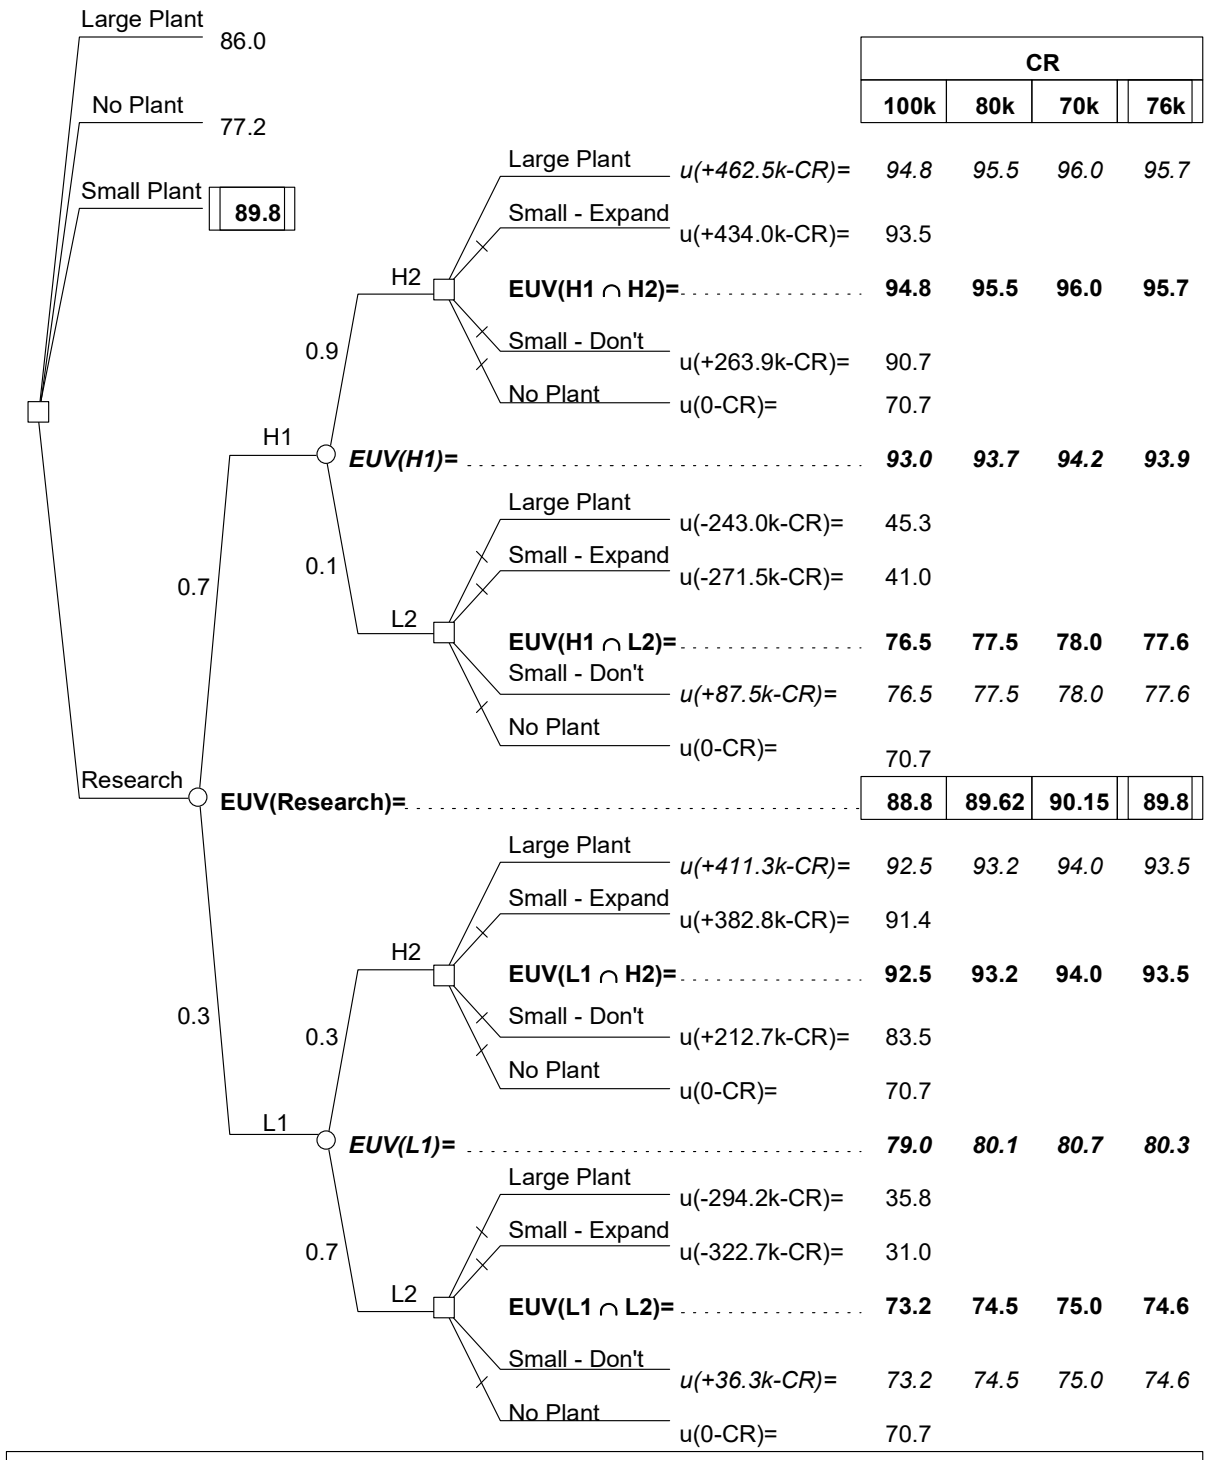
Example—Expected Value of Sampled (Imperfect) Information (EVSI).

Fig. E.26b shows the calculation of the *Expected Value of Sampled Information* (*EVSI*) for the source of aggregate example when using EMV. Fig. U.22 shows the EVSI for the source of aggregate example when using EUV and the utility curve in Fig. U.21. Notice that the cost values in Fig. E.26b have been transformed to profit in Fig. U.22 by assuming that "profit before paying for aggregate" is \$1,100,000. This transformation from cost to profit was necessary because the utility curve of Fig. U.21 shows the utility of profit (and not the utility of cost).

In Fig. U.22, test costs of CT=\$30,000 and \$40,000 were tried first. These values bracket CT* and thus a value of CT*= \$31,000 was determined by interpolation. Notice that:

$$\text{EVSI}_{\text{EUV}} = \$31{,}000 \;>\; \text{EVSI}_{\text{EMV}} = \$21{,}700$$

This illustrates that often a risk-averse decision maker is willing to pay more for the same test that reduces uncertainty than a risk-neutral decision maker.



Expected Value of Perfect Information (EVPI) -- Risk Averse Case

EVPI = CR such that:
Expected Utility given Perfect Information = Max Expected Utility without Perfect Information
For the case above:
EVPI = 80k-(80k-70k)(89.8-89.62)/(90.15-89.62) = 80k - 4k = $76k

Figure U.11 EVPI about Demand Over Ten Years (Risk-Averse Behavior)

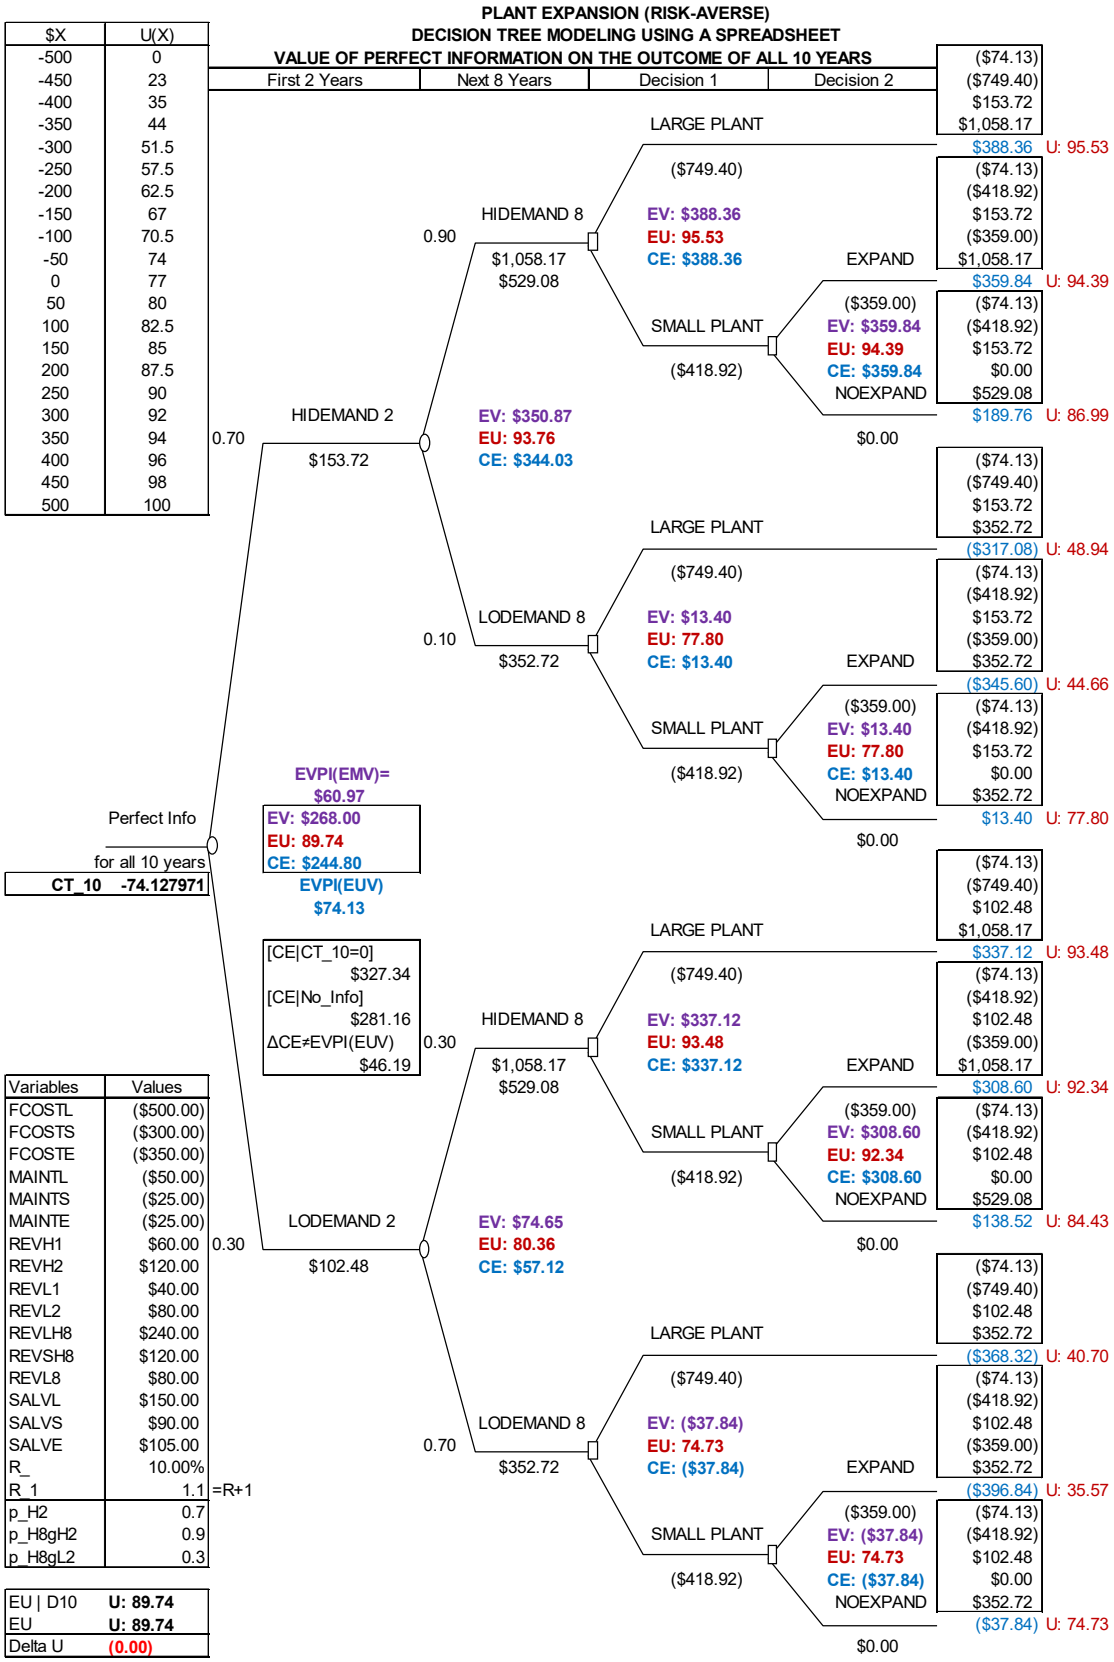


Figure U.12 EVPI about Demand Over Ten Years (Excel) (Risk-Averse Behavior)

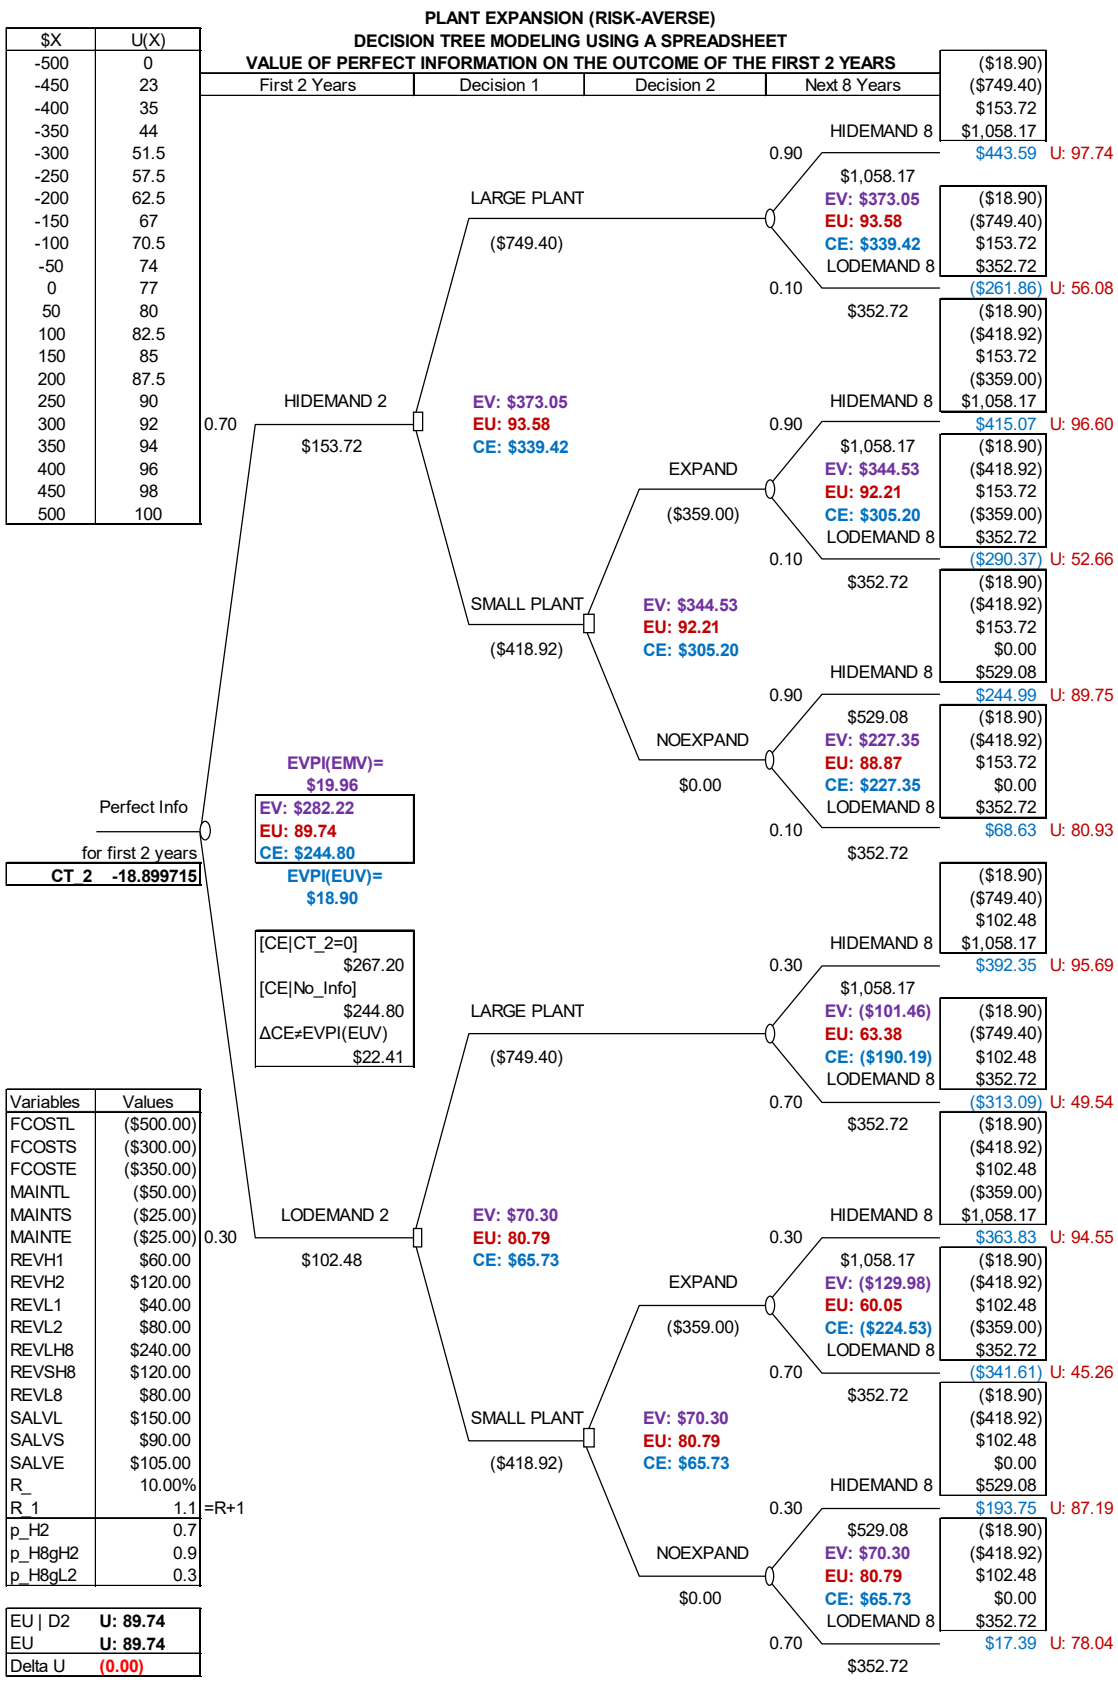


Figure U.13 EVPI about Demand Over First Two Years (Excel) (Risk-Averse Behavior)

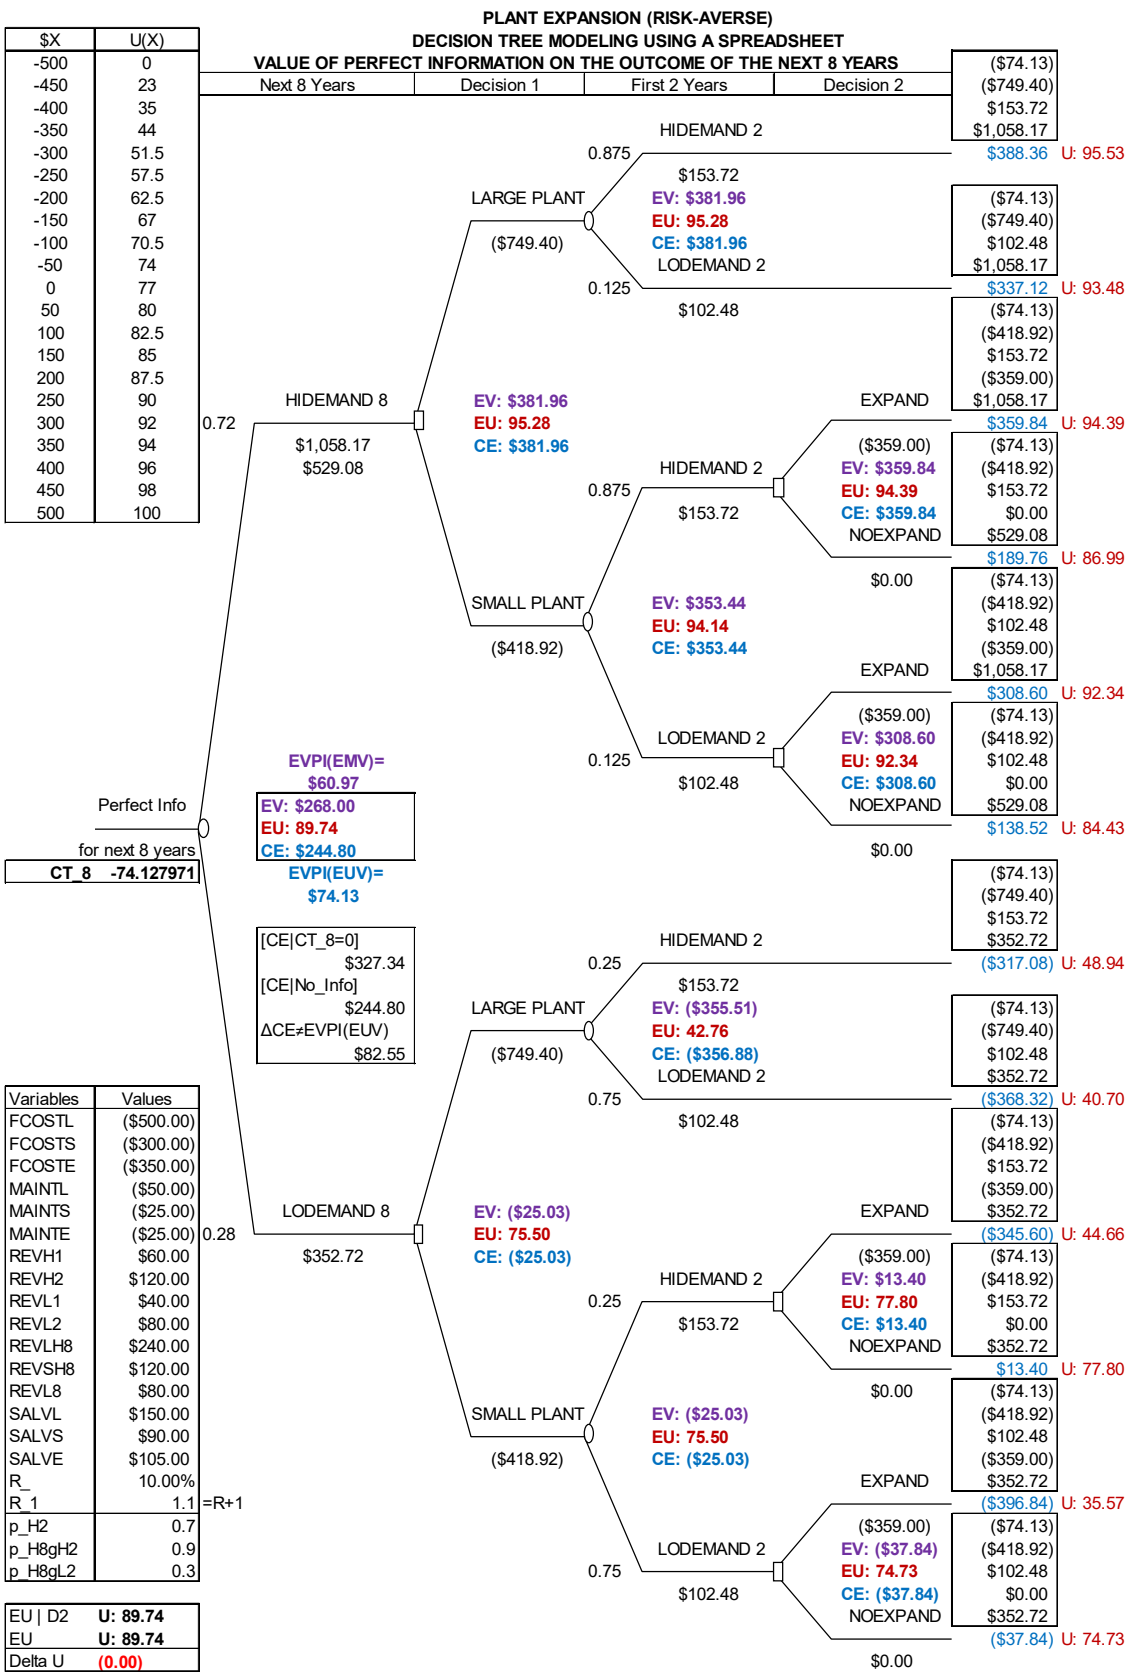


Figure U.14 EVPI about Demand Over Next Eight Years (Excel) (Risk-Averse Behavior)

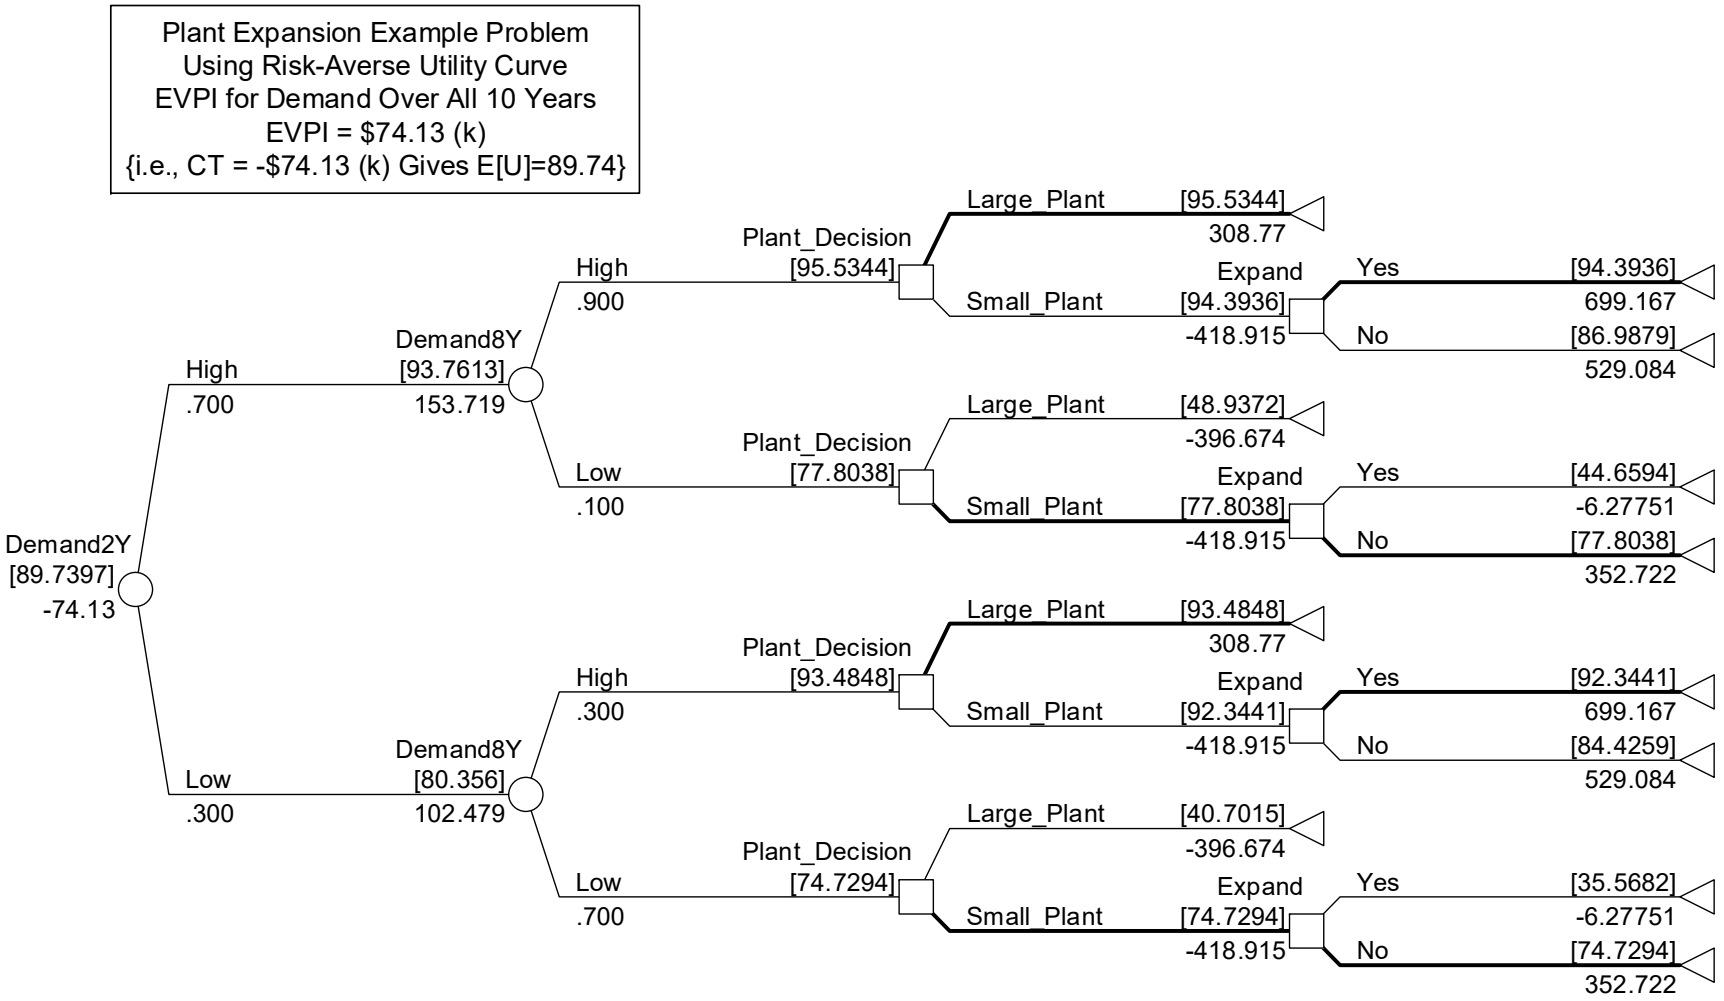


Figure U.15 EVPI about Demand Over Ten Years (DPL using EU)

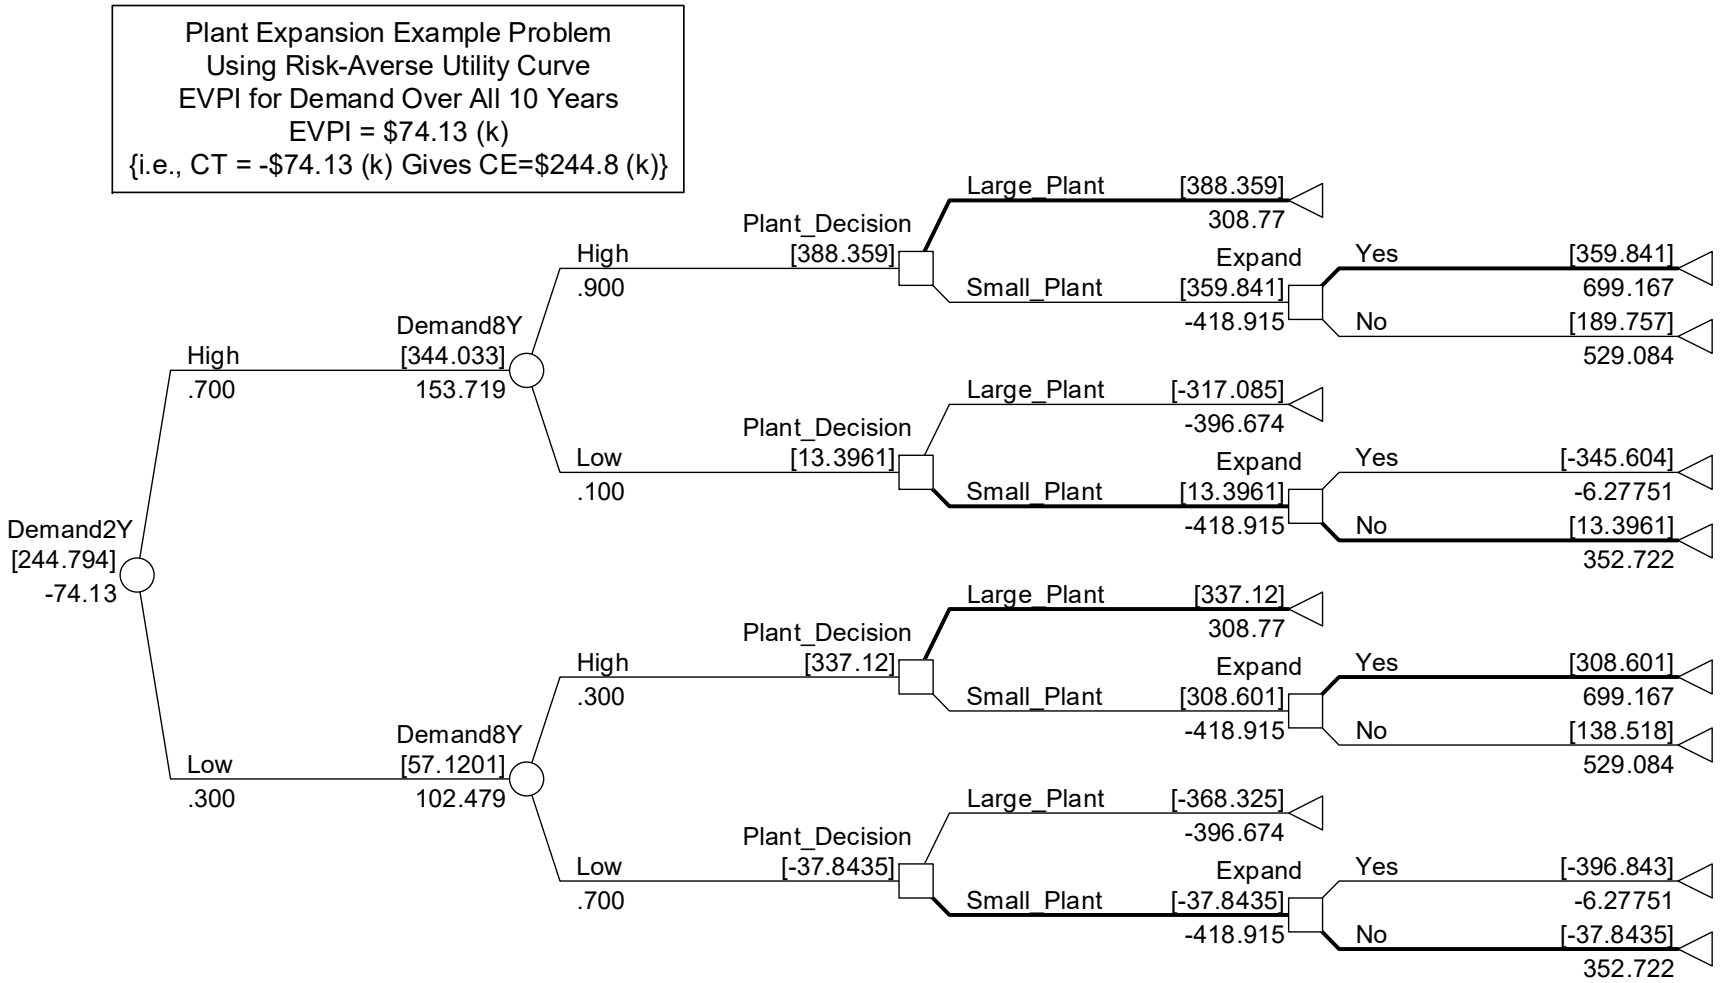


Figure U.16 EVPI about Demand over Ten Years (DPL using CE)

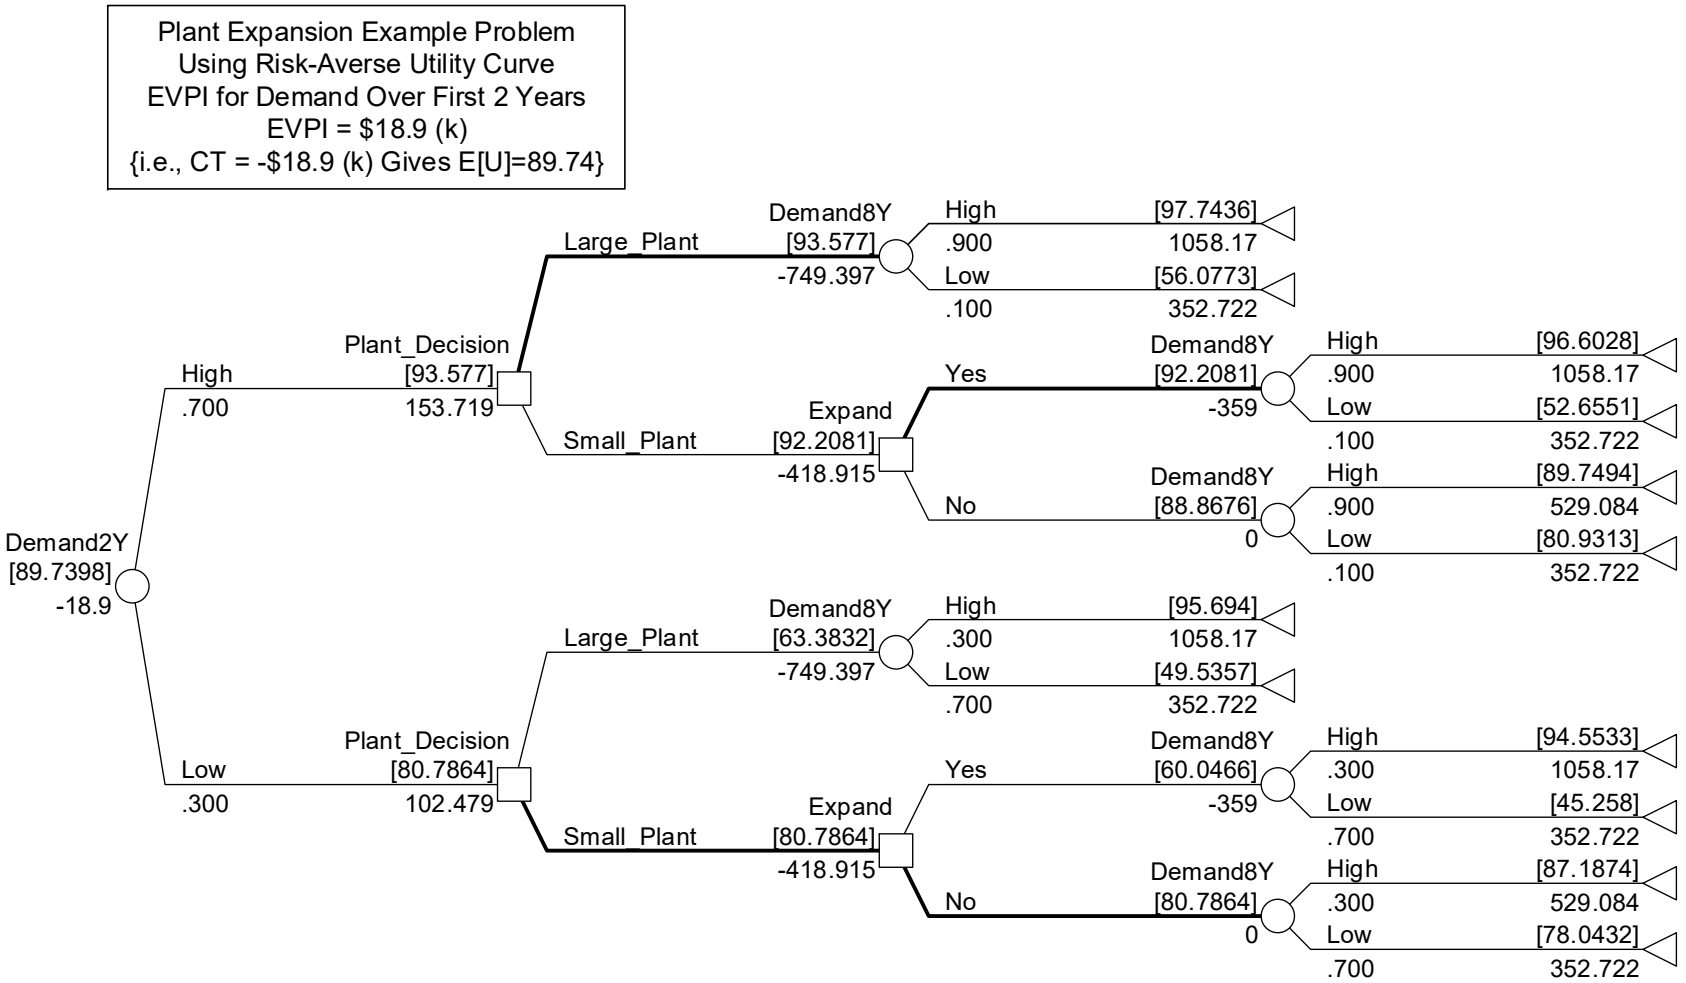


Figure U.17 EVPI about Demand over First Two Years (DPL using EU)

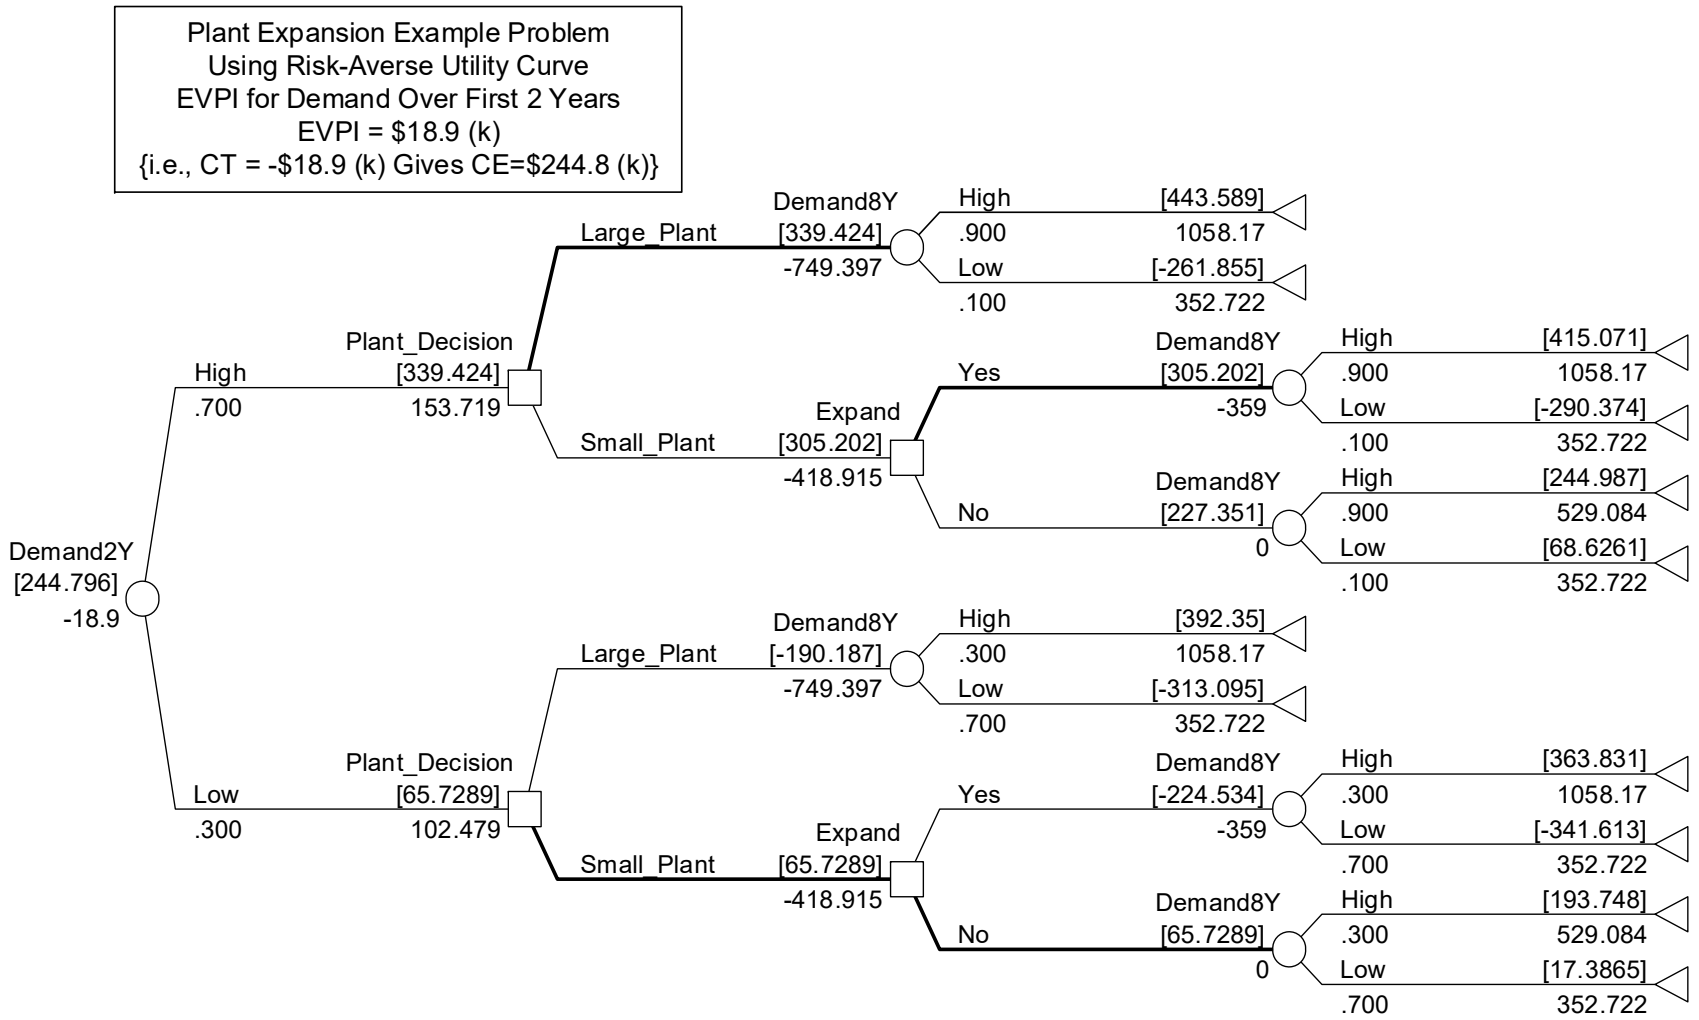


Figure U.18 EVPI about Demand over First Two Years (DPL using CE)

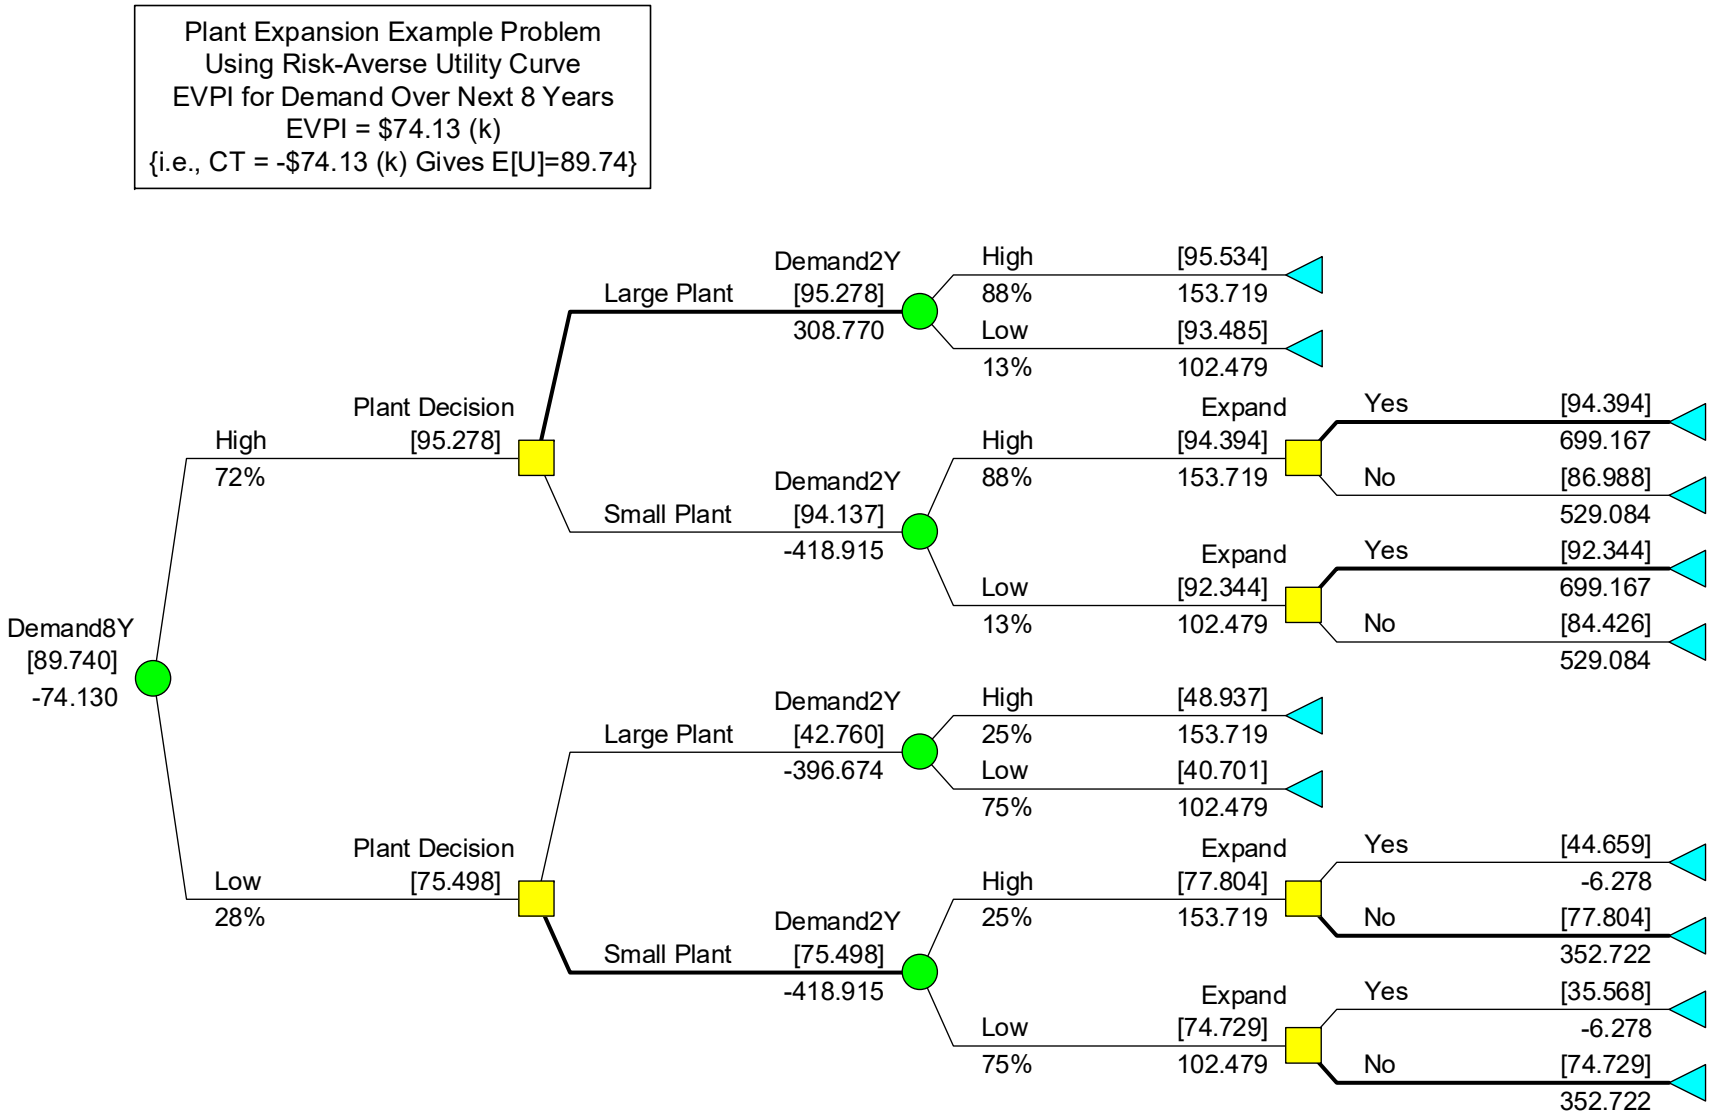


Figure U.19 EVPI about Demand over Next Eight Years (DPL using EU)

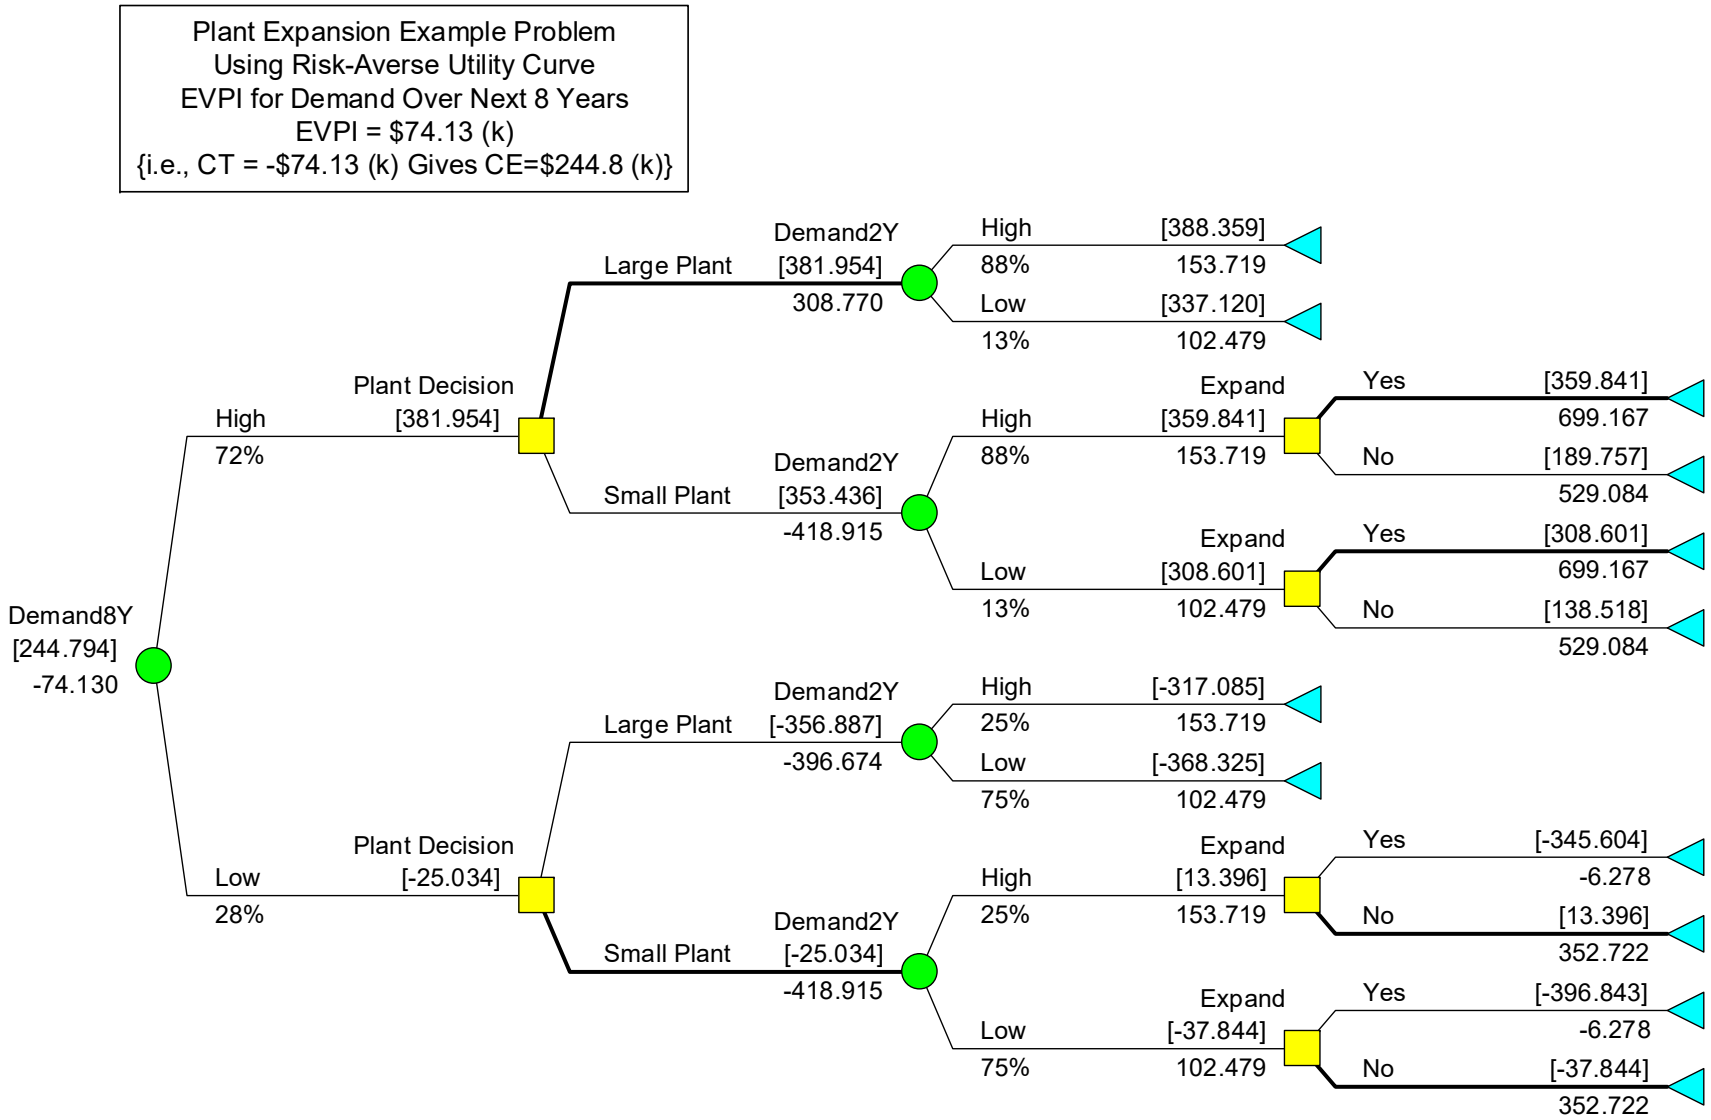
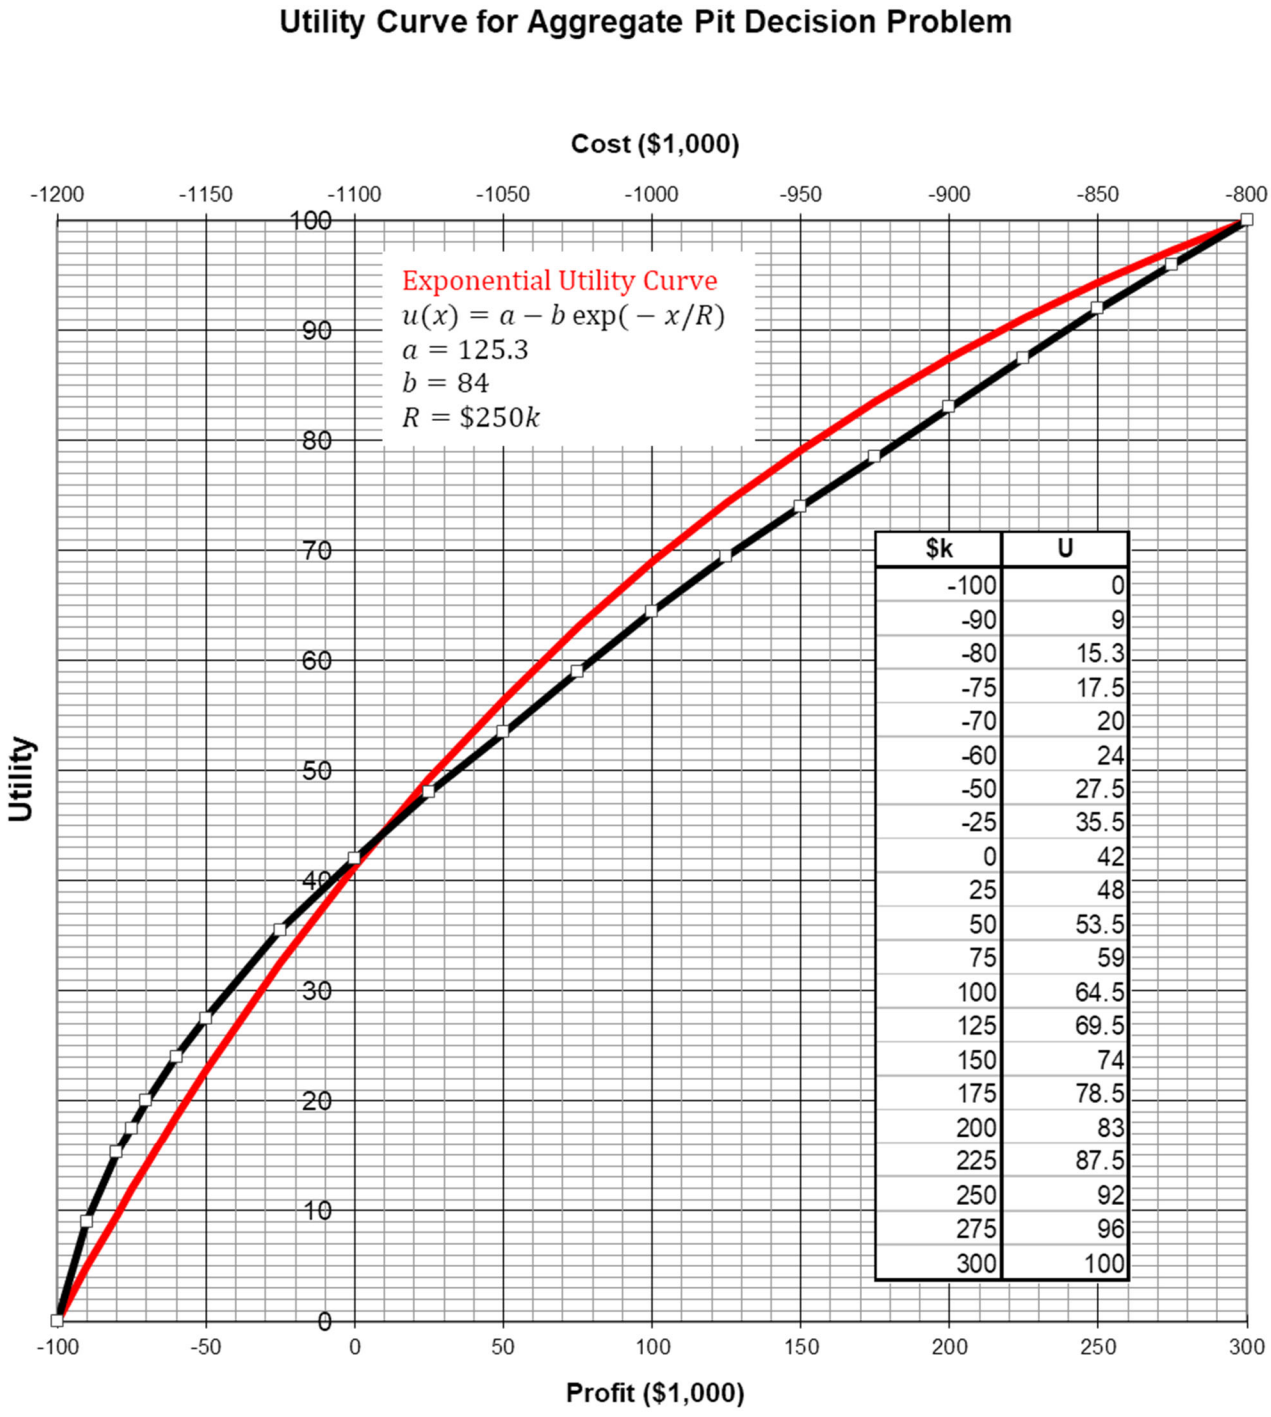
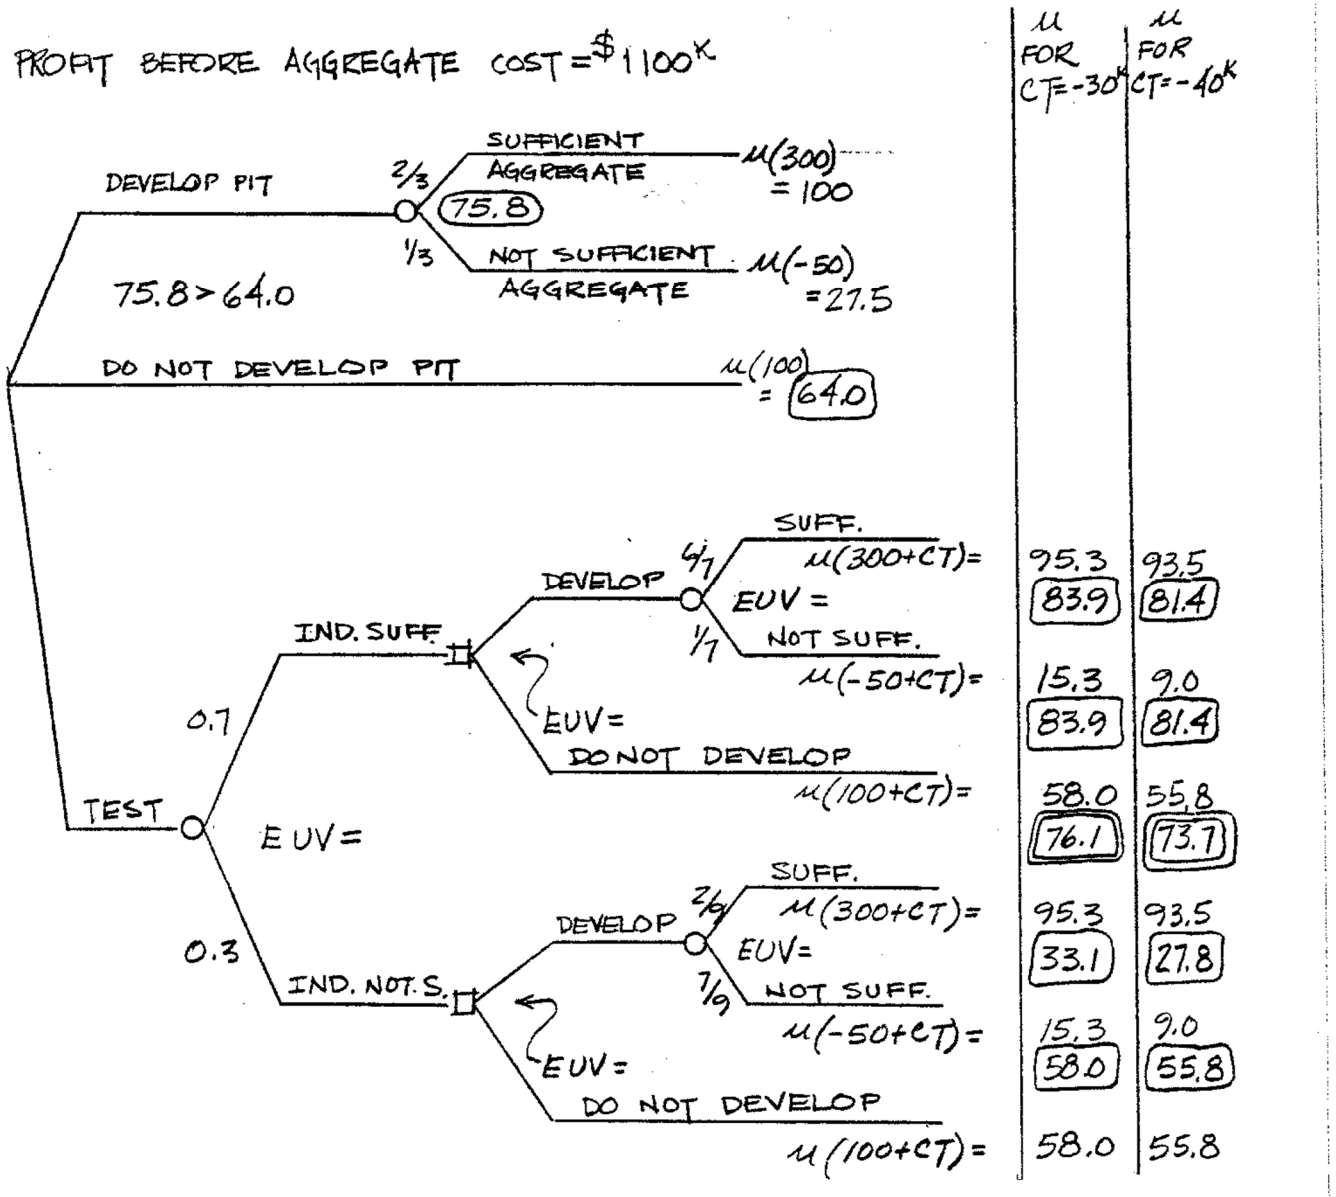


Figure U.20 EVPI about Demand over Next Eight Years (DPL using CE)



$k	U
-100	0
-90	9
-80	15.3
-75	17.5
-70	20
-60	24
-50	27.5
-25	35.5
0	42
25	48
50	53.5
75	59
100	64.5
125	69.5
150	74
175	78.5
200	83
225	87.5
250	92
275	96
300	100

Figure U.21 Utility Function for Decision to Develop Aggregate Pit



VALUE OF TEST = CT AT WHICH EUV (WITH COST OF TEST = CT) = EUV (W/O TEST)

= CT AT WHICH EUV (WITH CT) = 75.8

$$= 40^K - \frac{75.8 - 73.7}{76.1 - 73.7}(40-30) = 40^K - 8.8^K = \$31{,}000$$

Figure U.22. EVSI for Development of Aggregate Pit using Utility

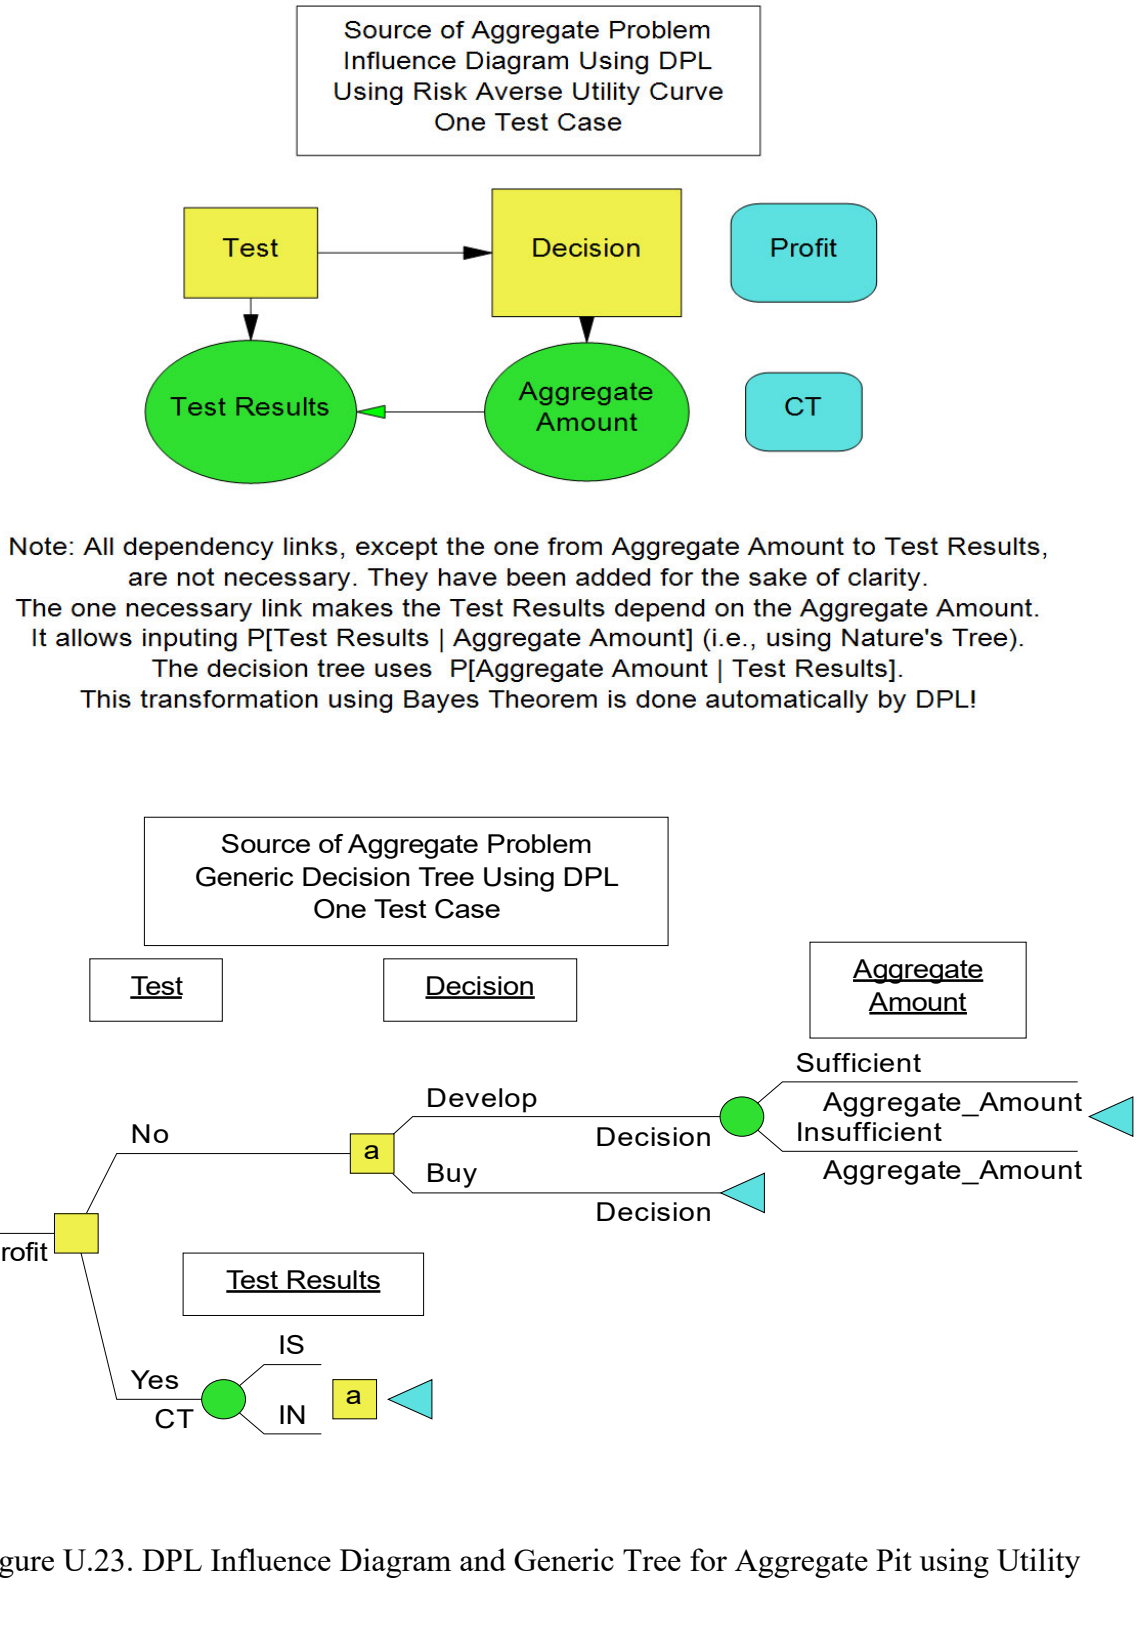


Figure U.23. DPL Influence Diagram and Generic Tree for Aggregate Pit using Utility

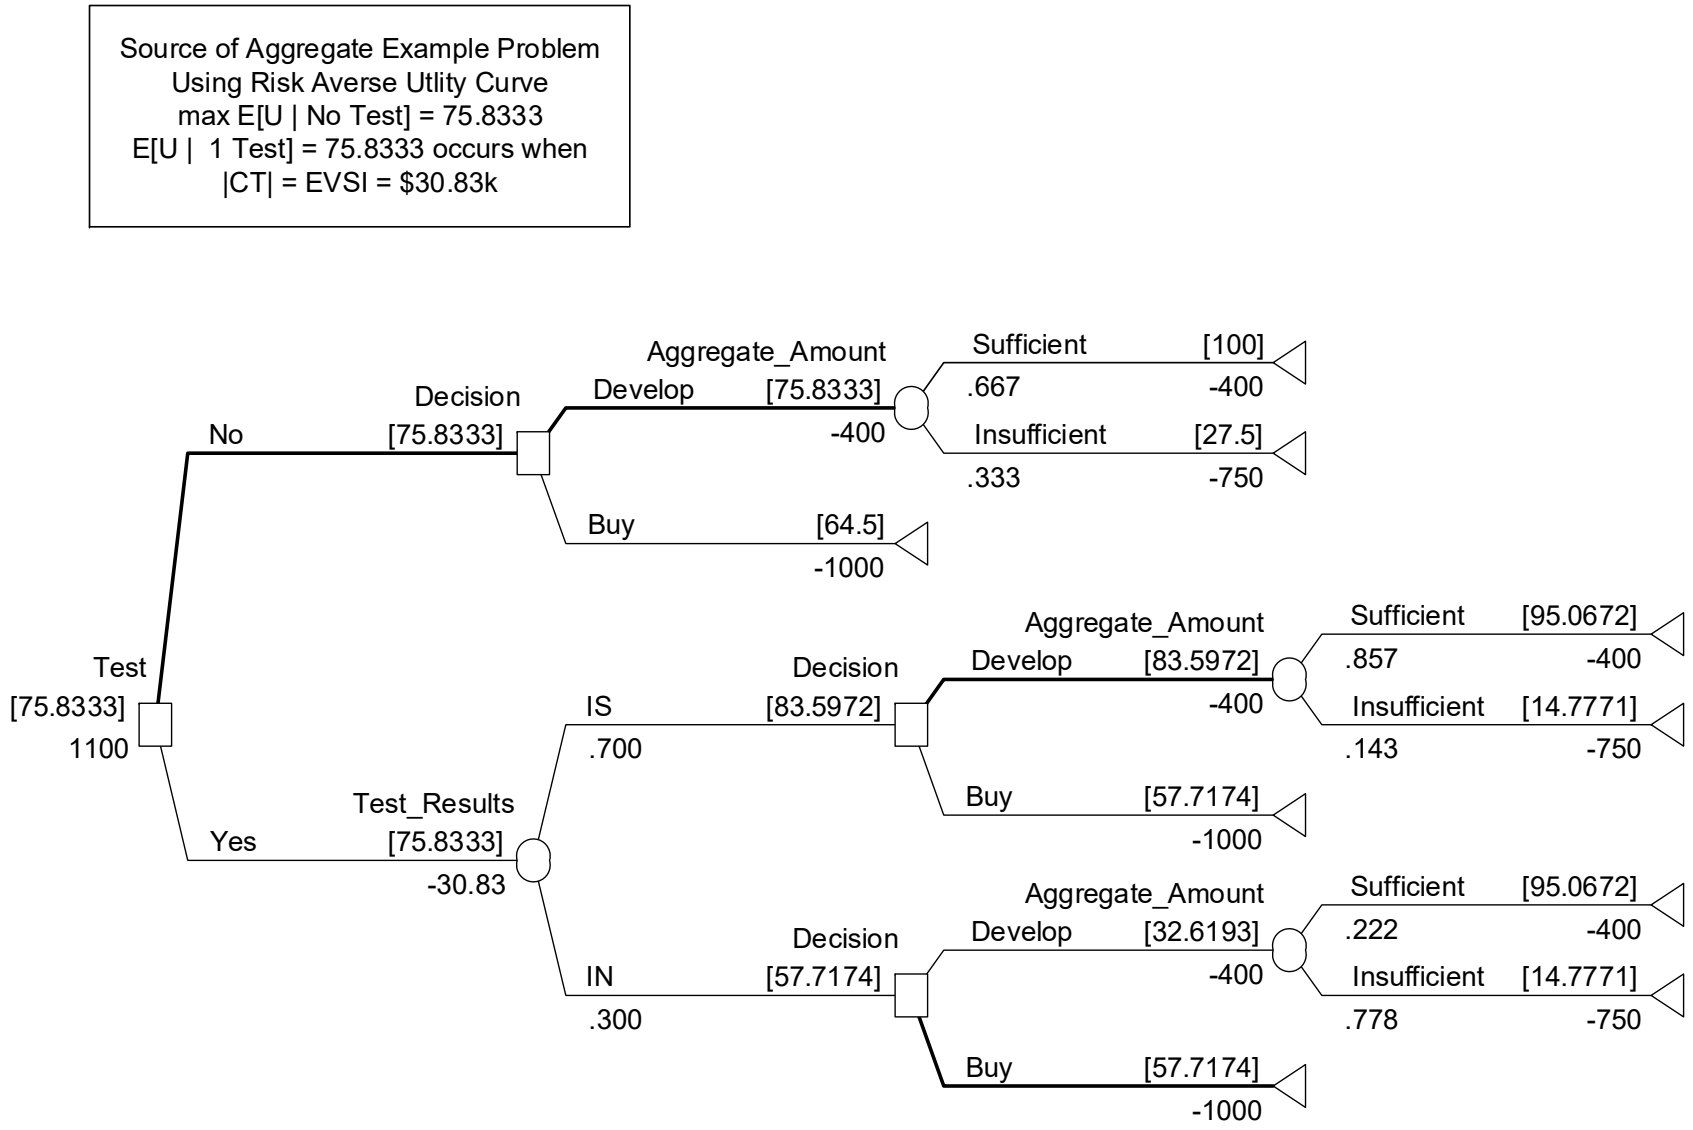


Figure U.24. EVSI for Aggregate Amount in Pit (DPL using EU)

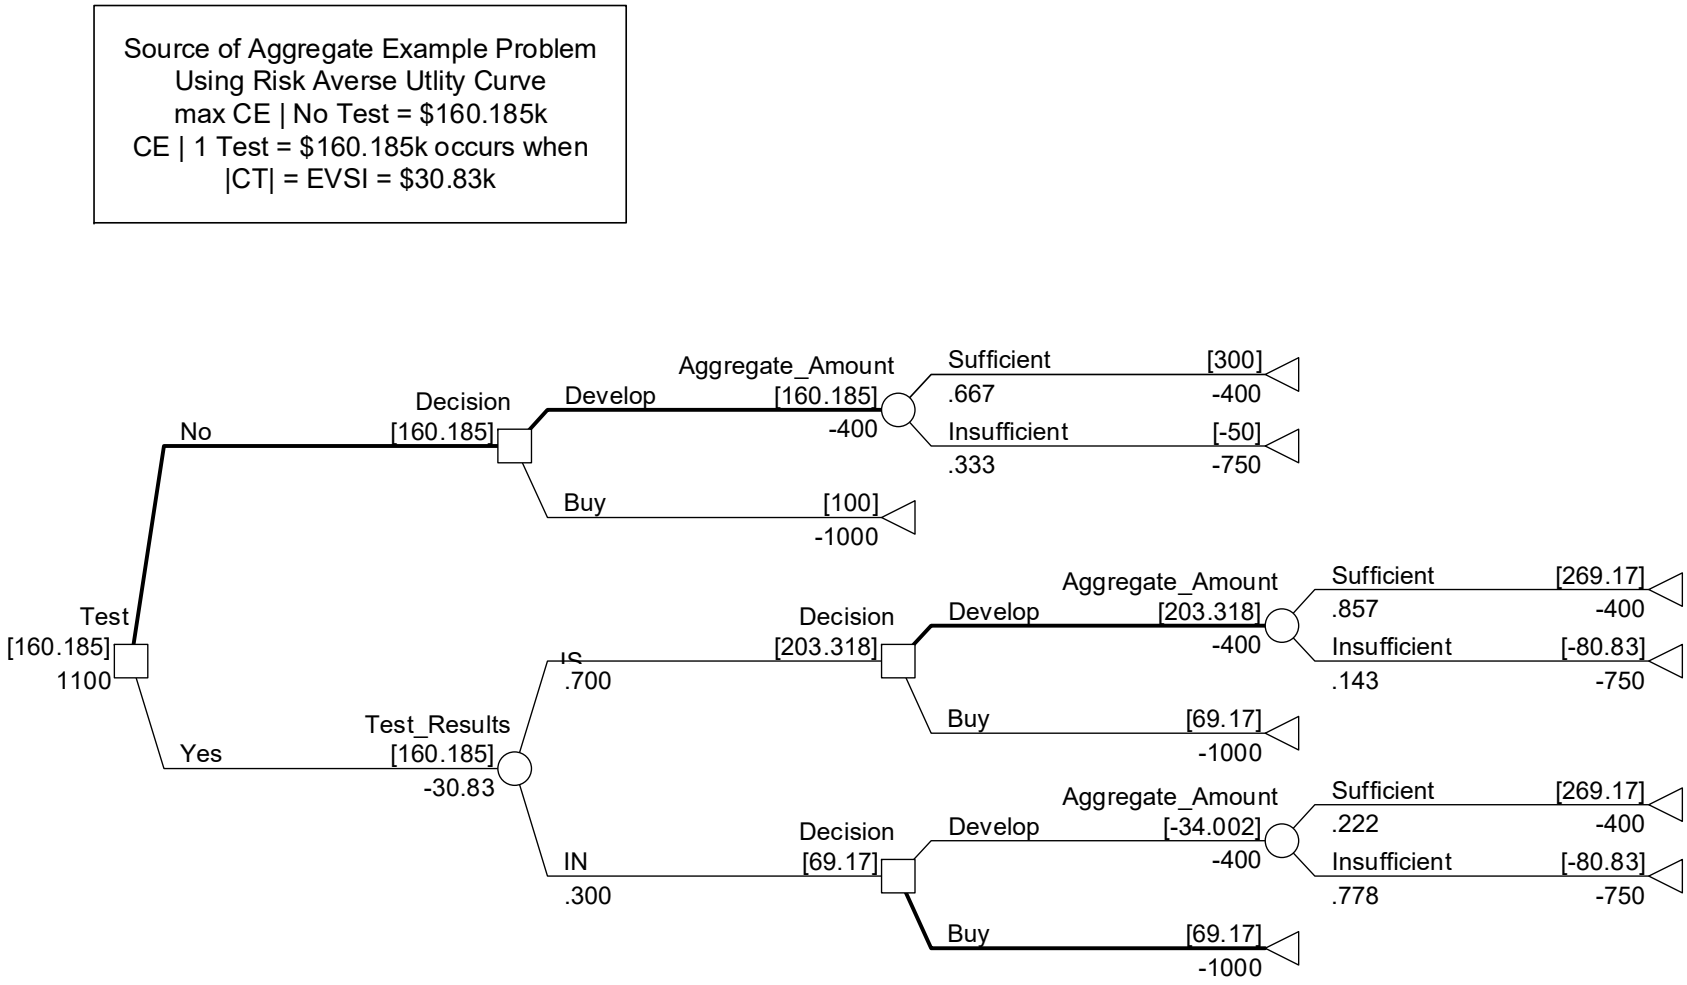


Figure U.25. EVSI for Aggregate Amount in Pit (DPL using CE)

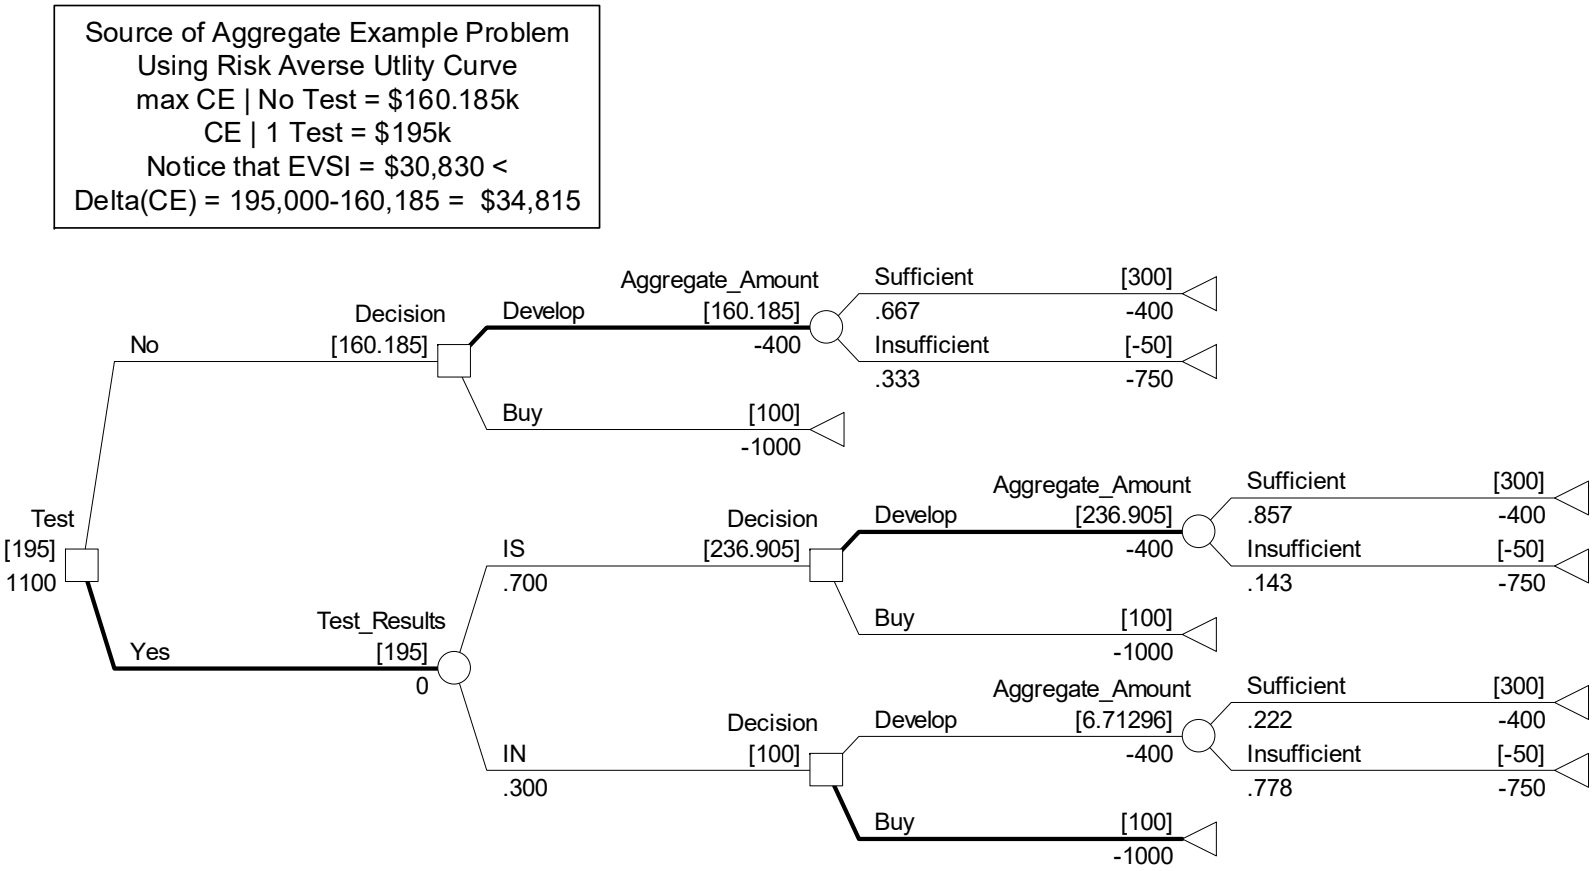


Figure U.26. CE for No-Cost Test about Aggregate Amount in Pit (DPL)

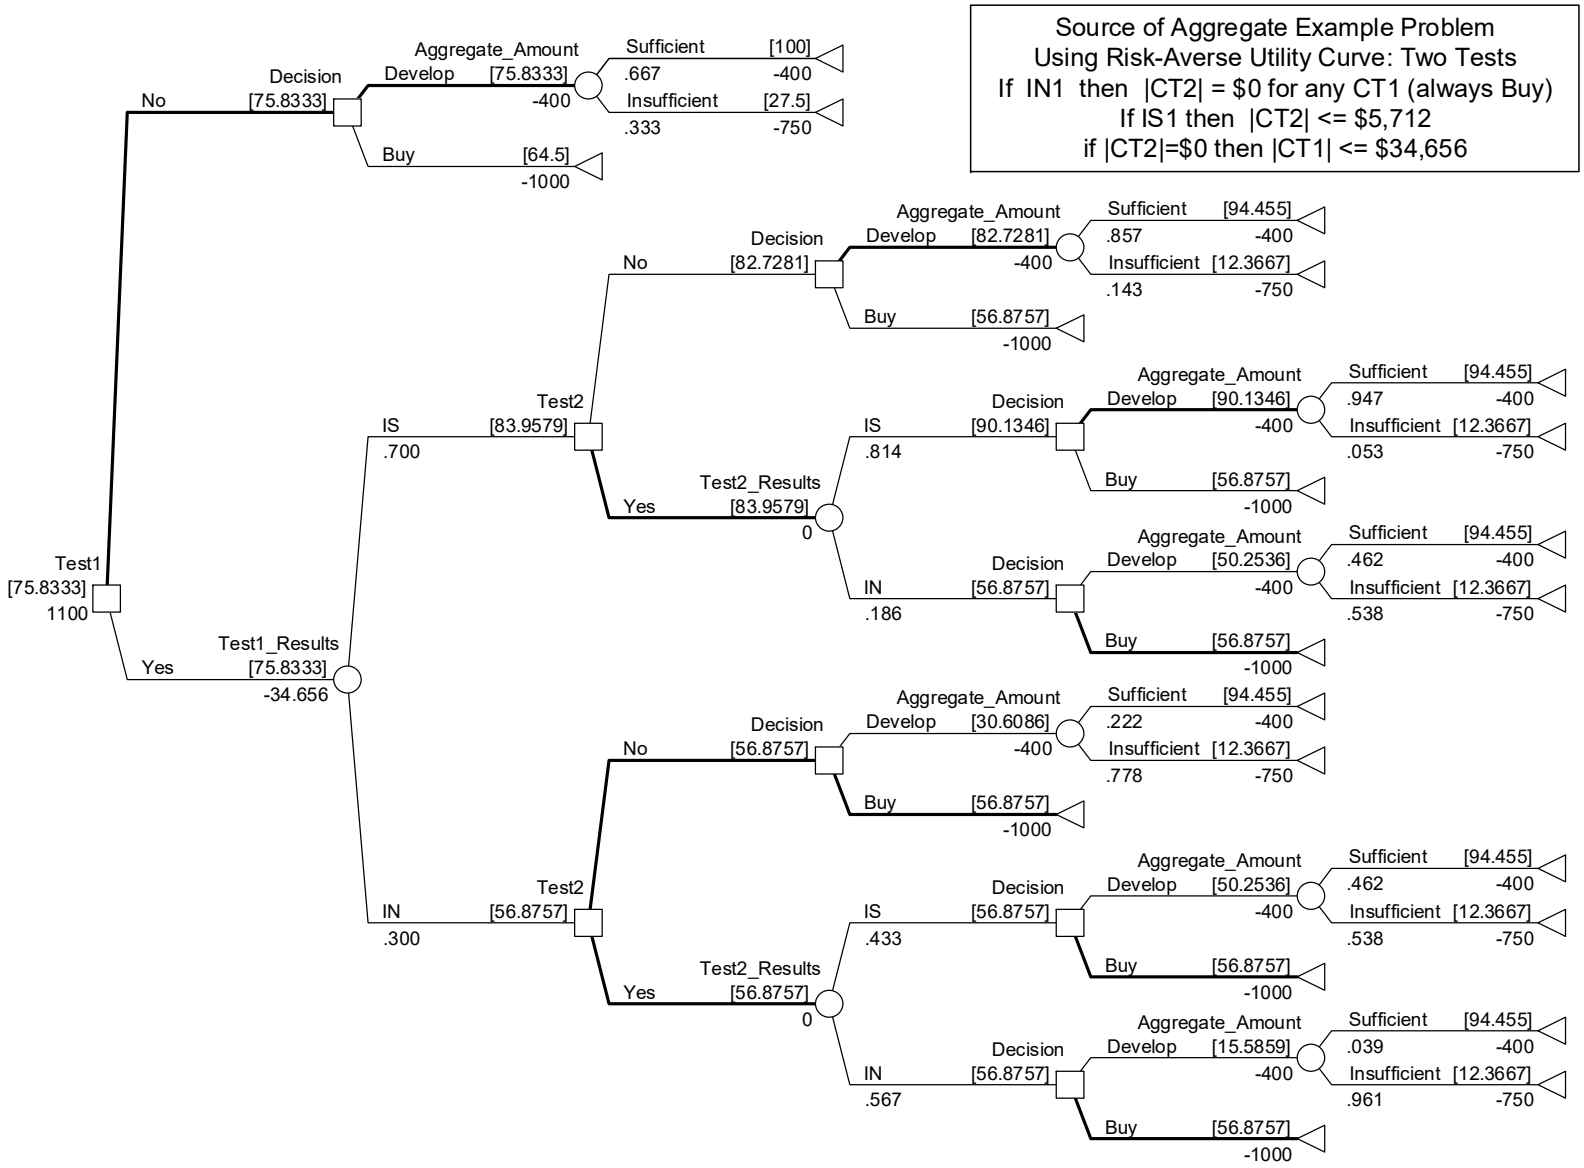


Figure U.27. EVSI for 2 Tests about Aggregate Amount in Pit (DPL using CT2=0)

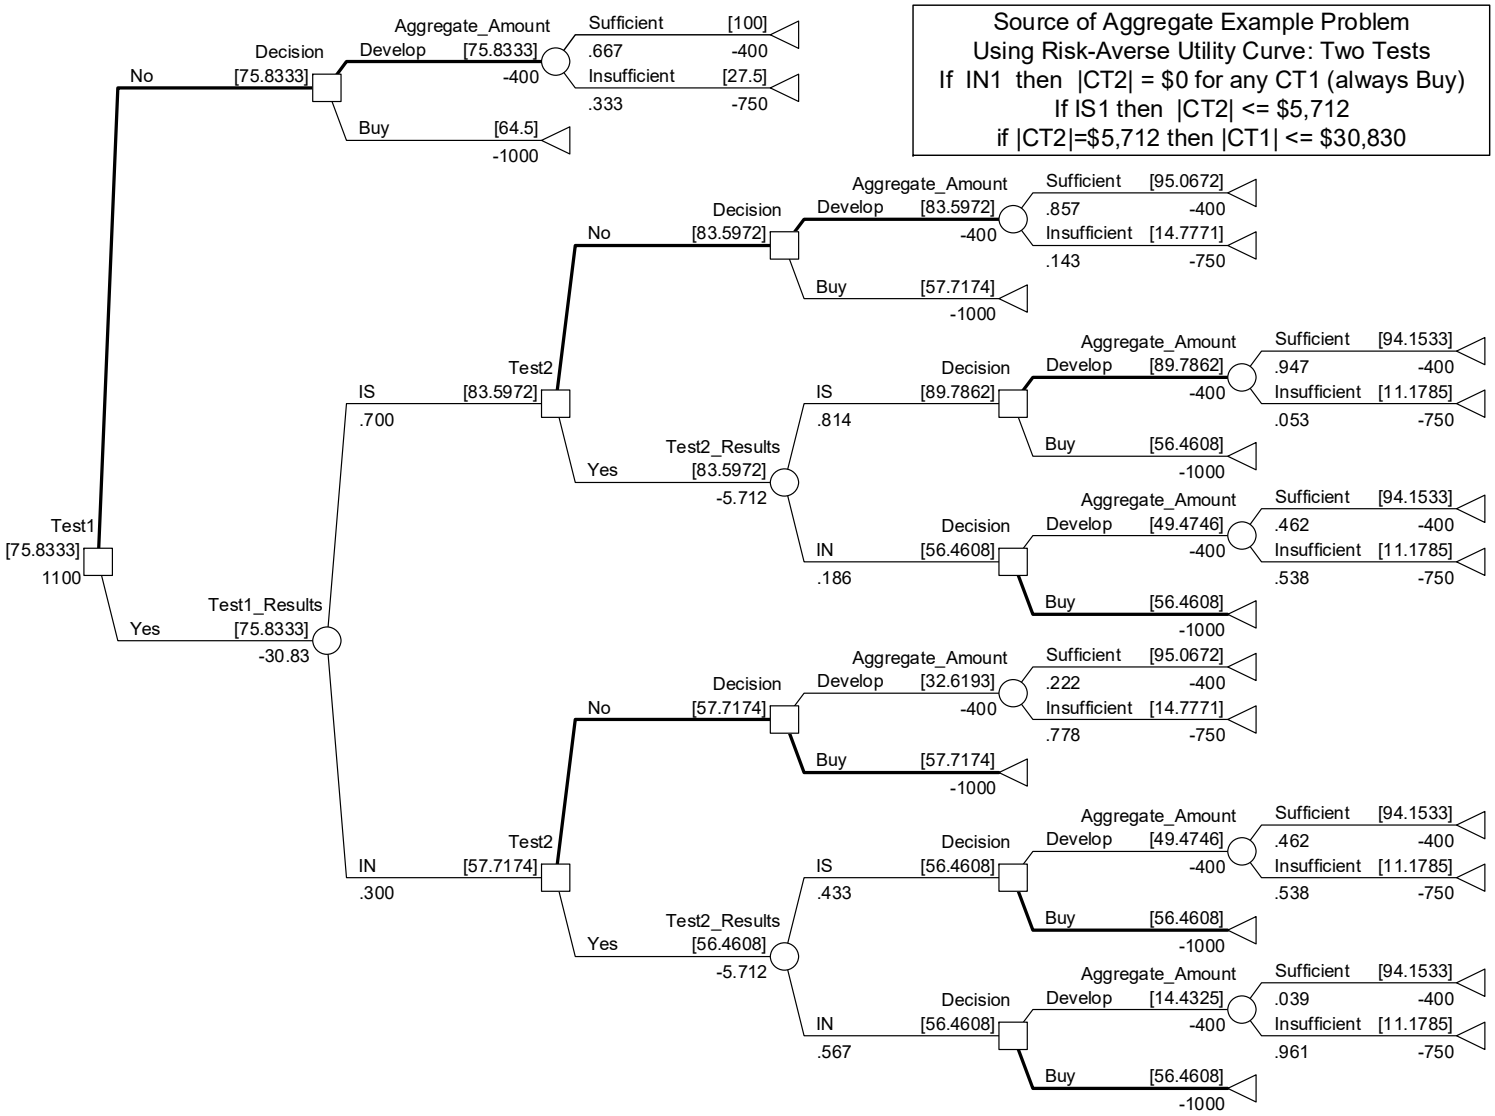


Figure U.28. EVSI for 2 Tests about Aggregate Amount in Pit (DPL using CT2=$5,712)

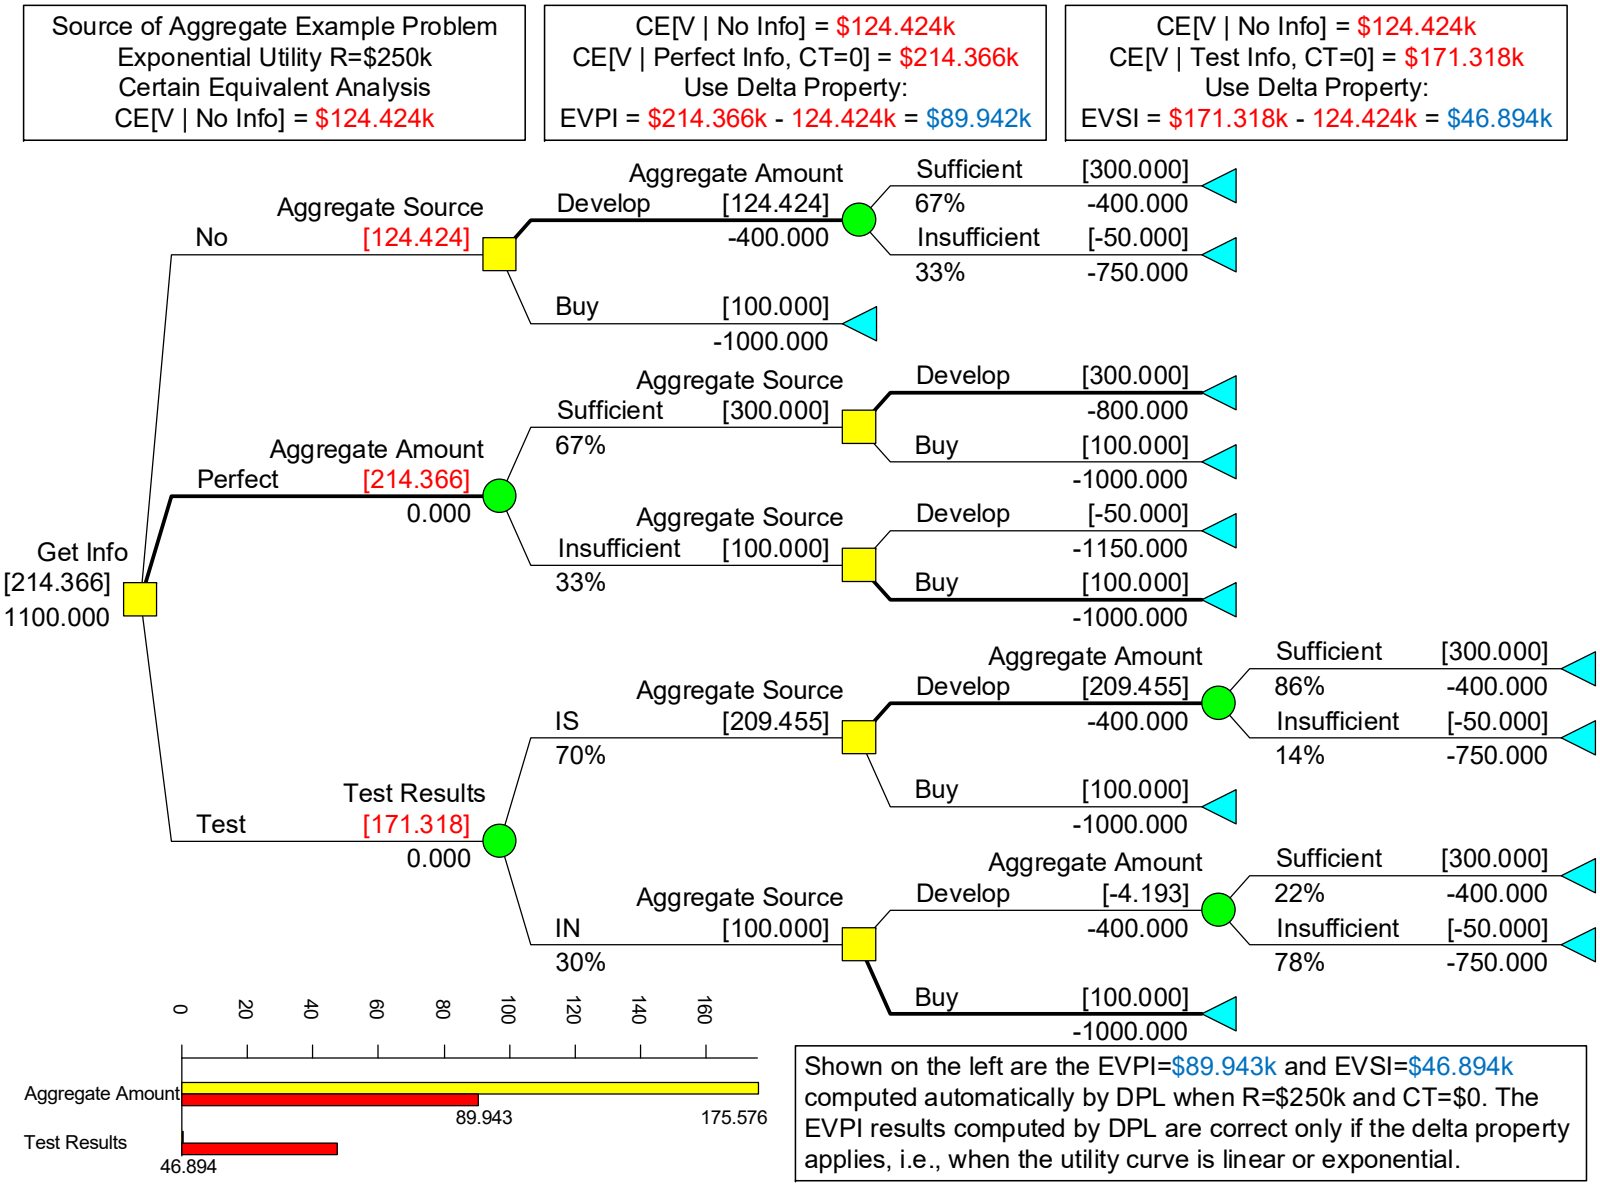


Figure U.29. 1-Test EVPI & EVSI for Aggregate Amount using Exponential Utility, R=$250k

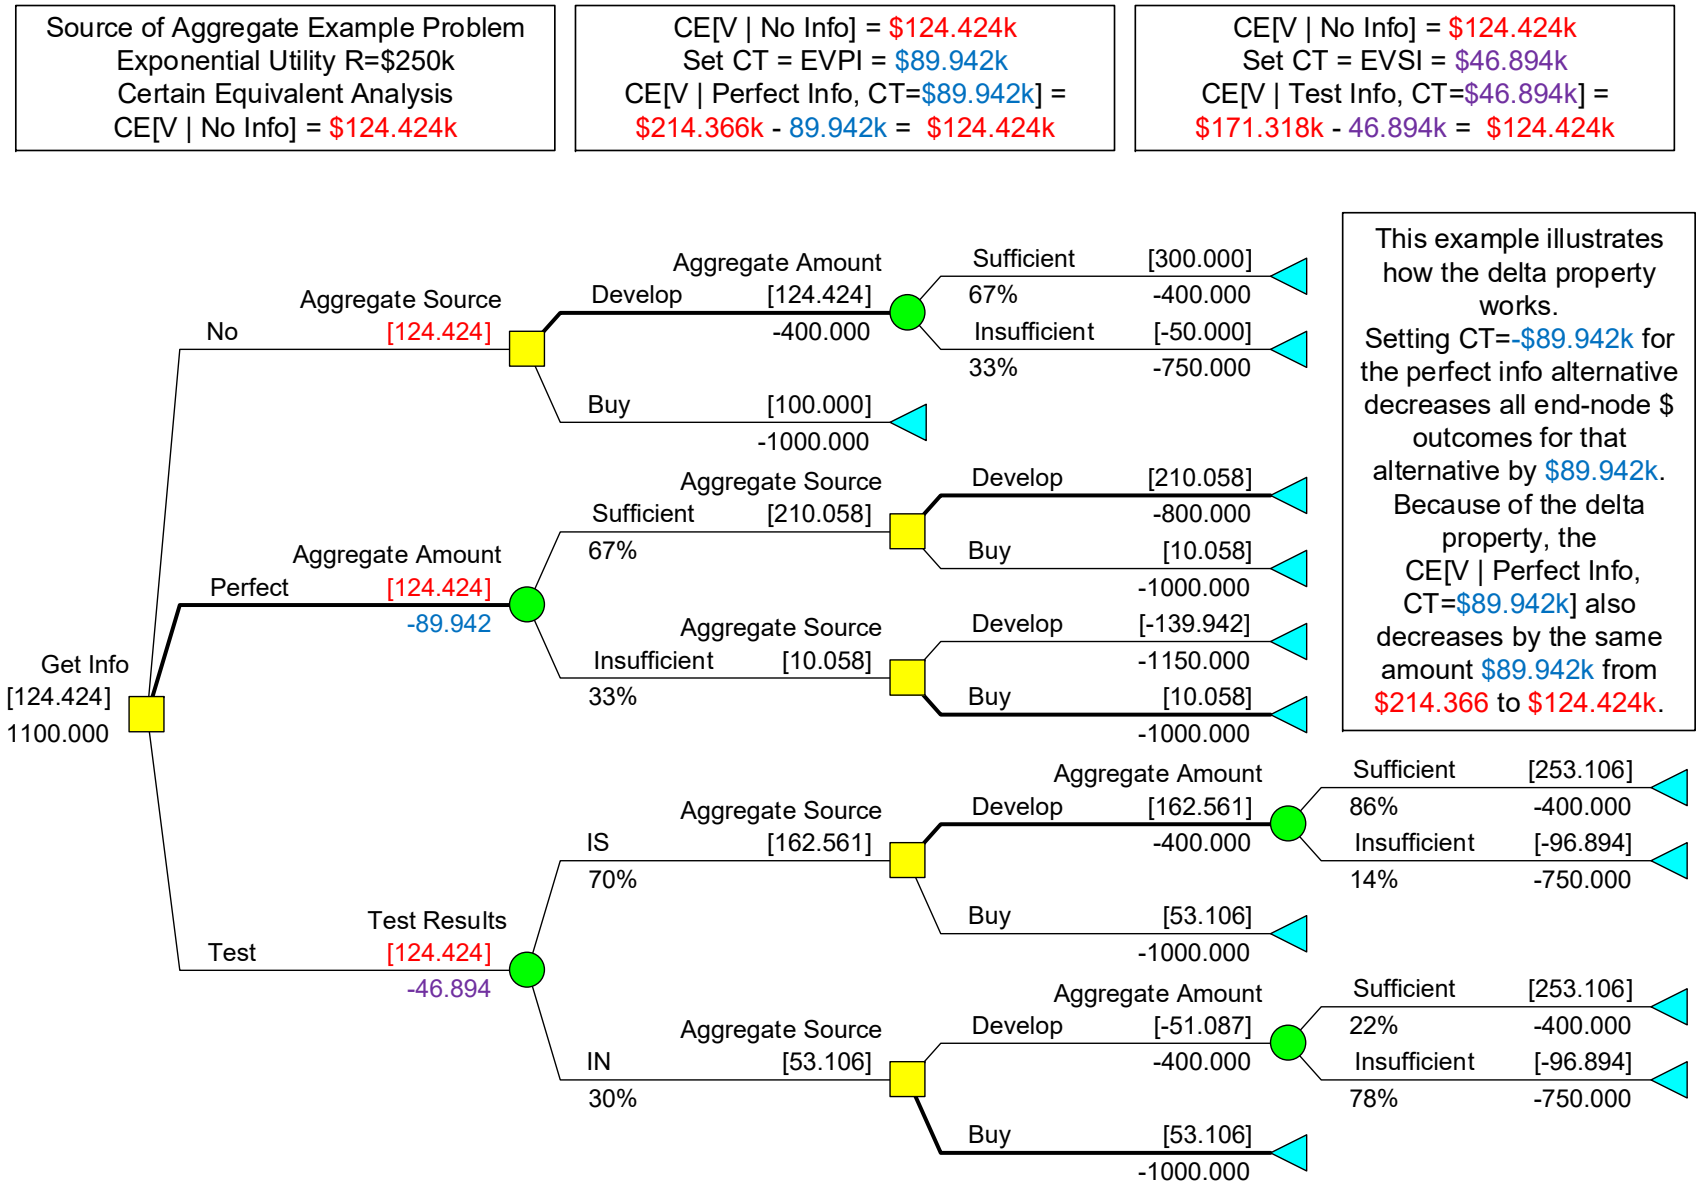


Figure U.30. CT at EVPI & EVSI for Aggregate Amount using Exponential Utility, R=$250k

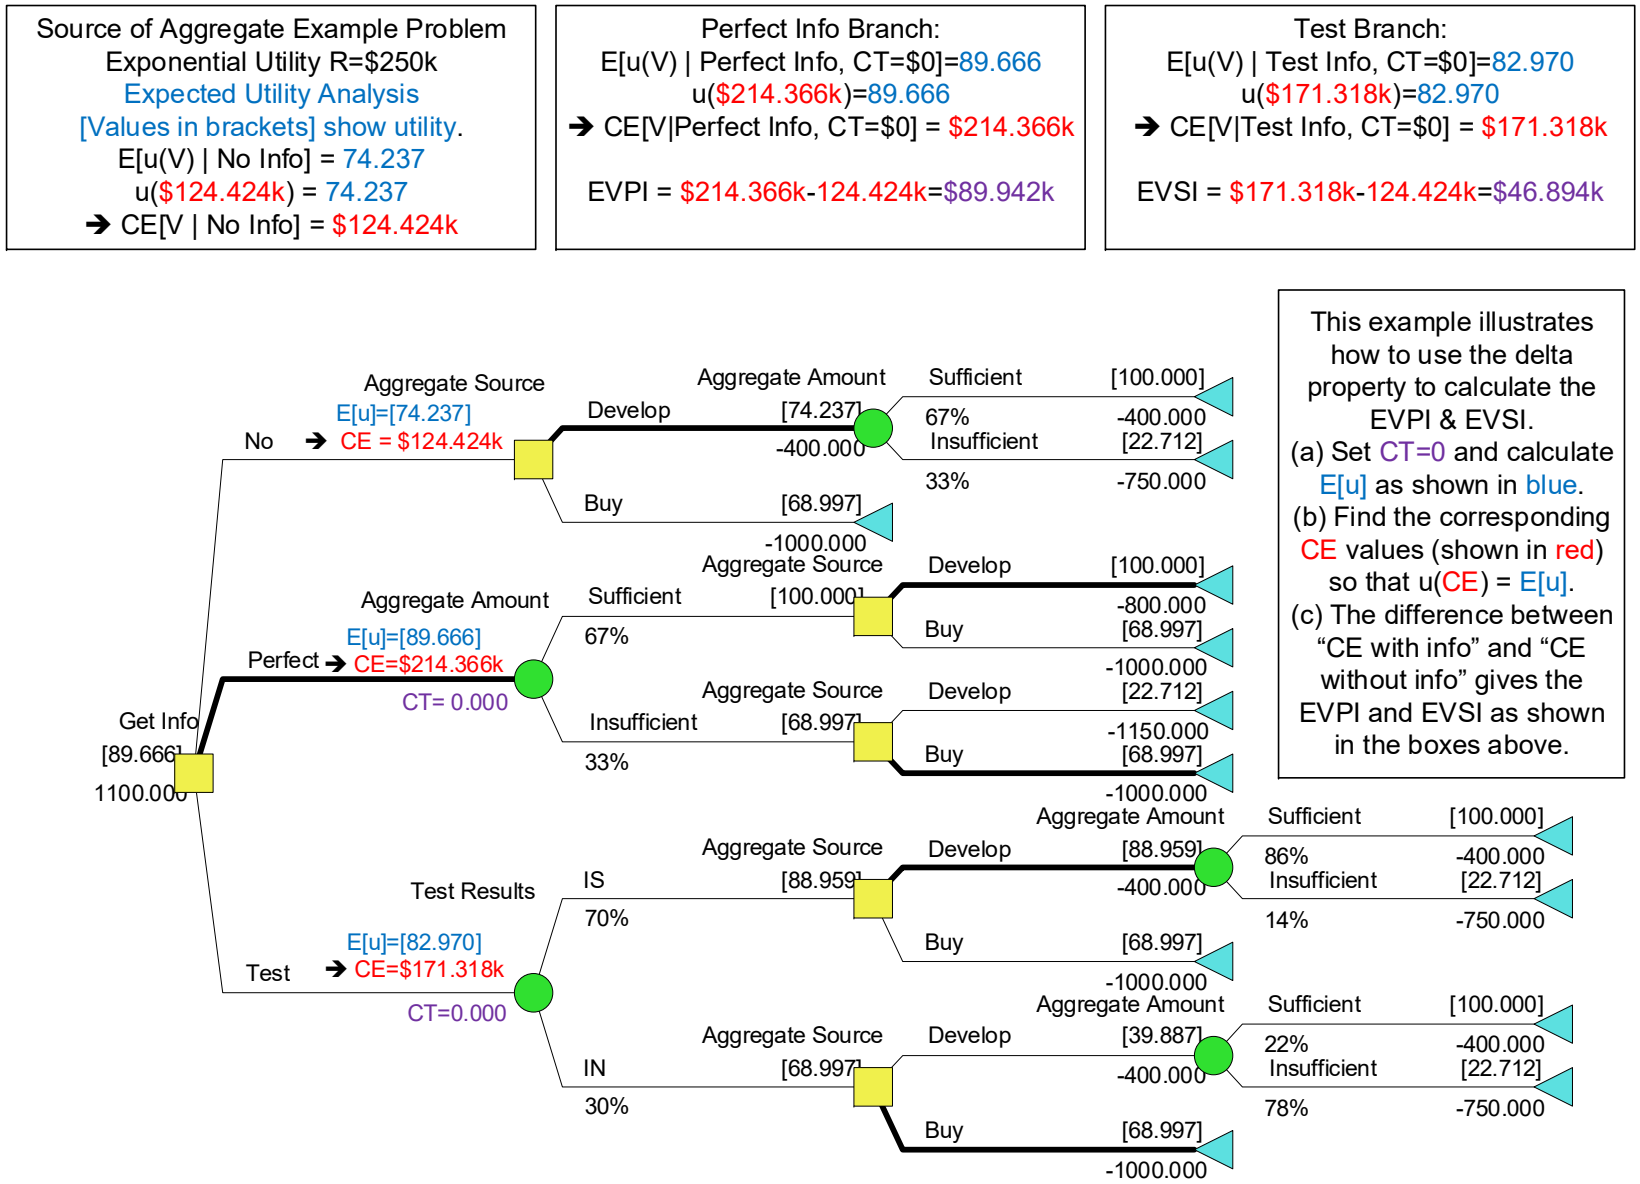


Figure U.31. EVPI & EVSI for Aggregate Amount using Exponential Utility, R=$250k

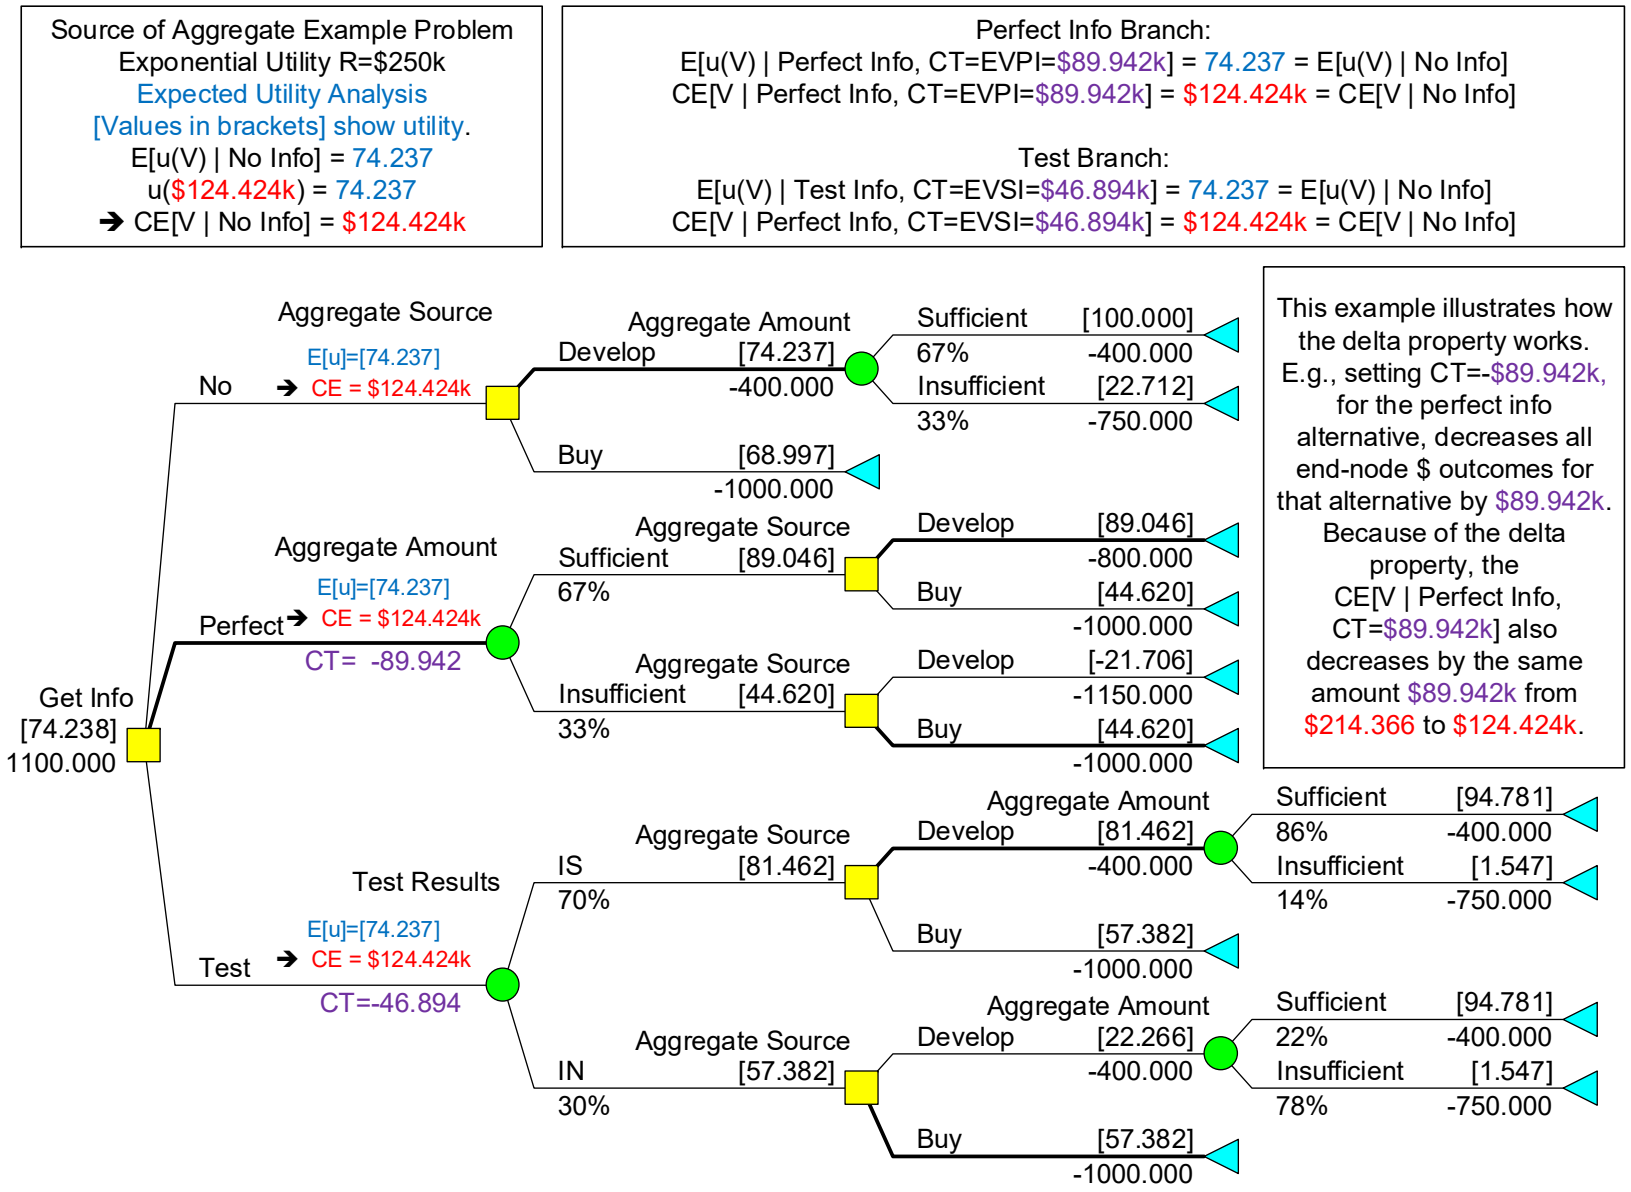


Figure U.32. CT at EVPI & EVSI for Aggregate Amount using Exponential Utility, R=$250k

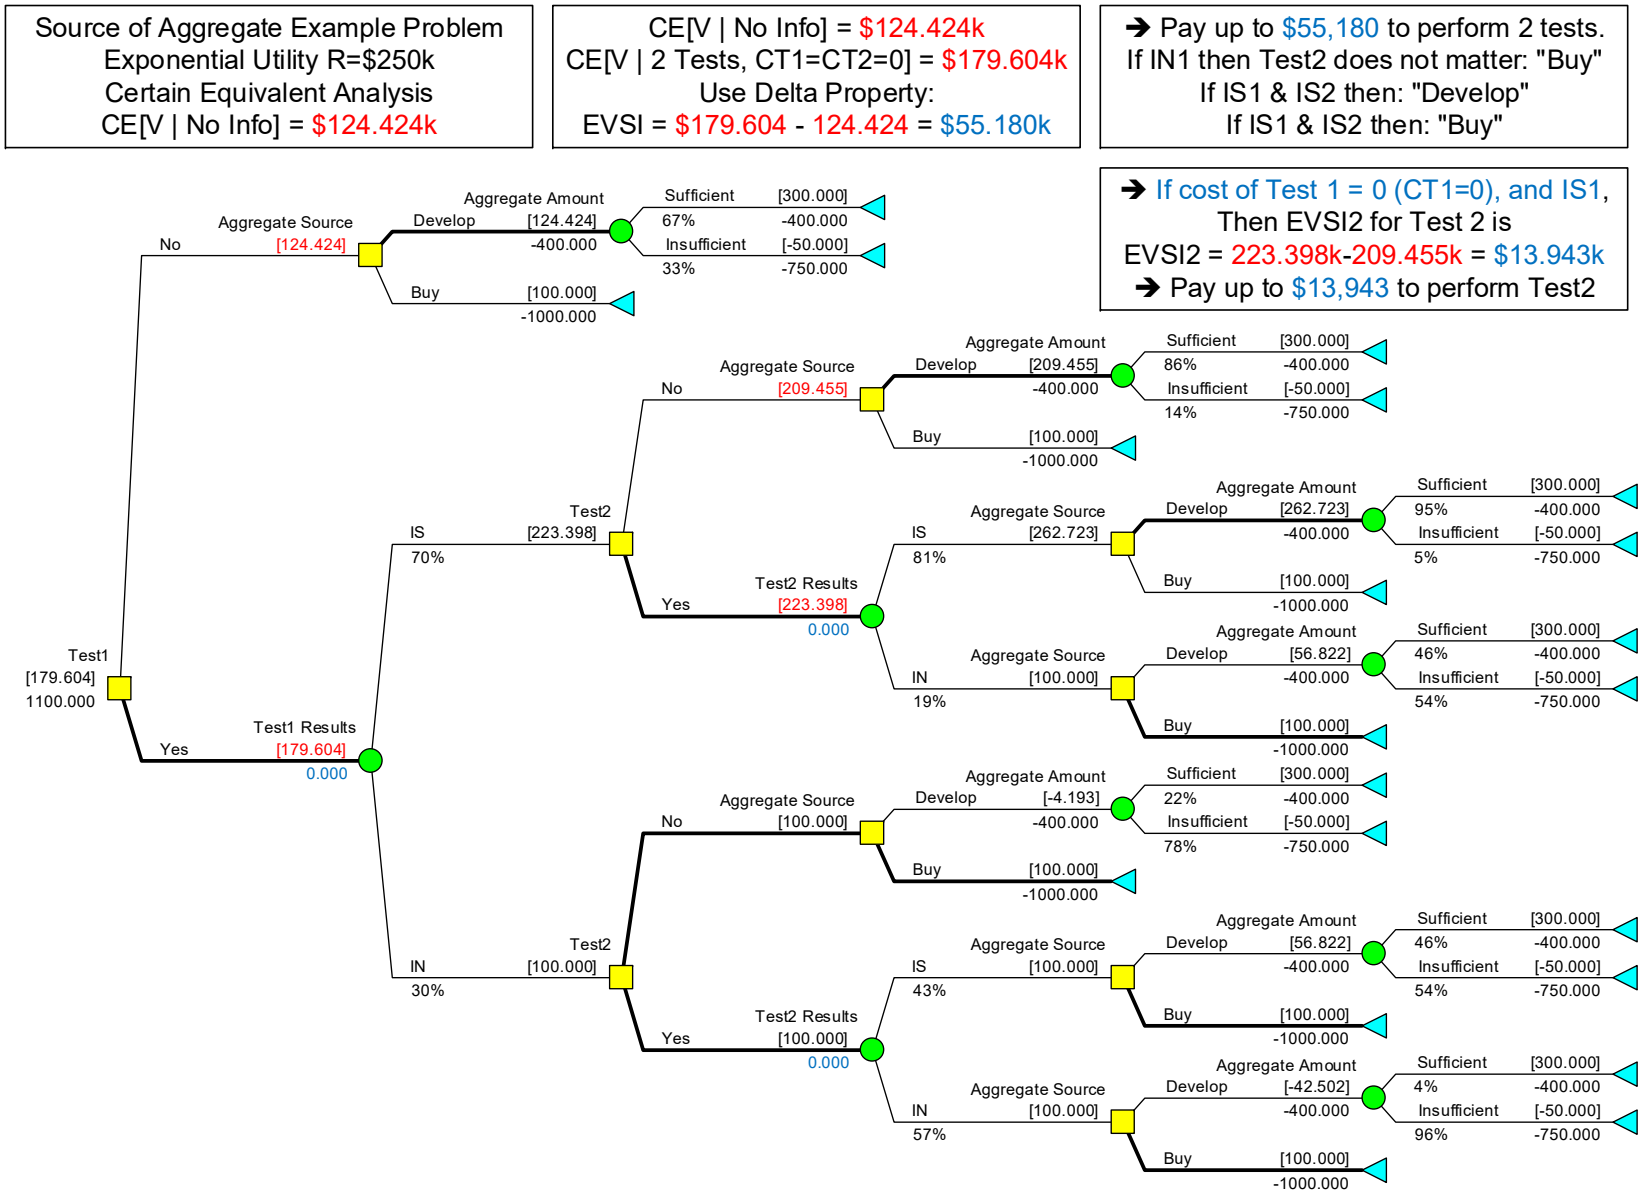


Figure U.33. EVSI for 2 Tests on Aggregate Amount using Exponential Utility, R=$250k

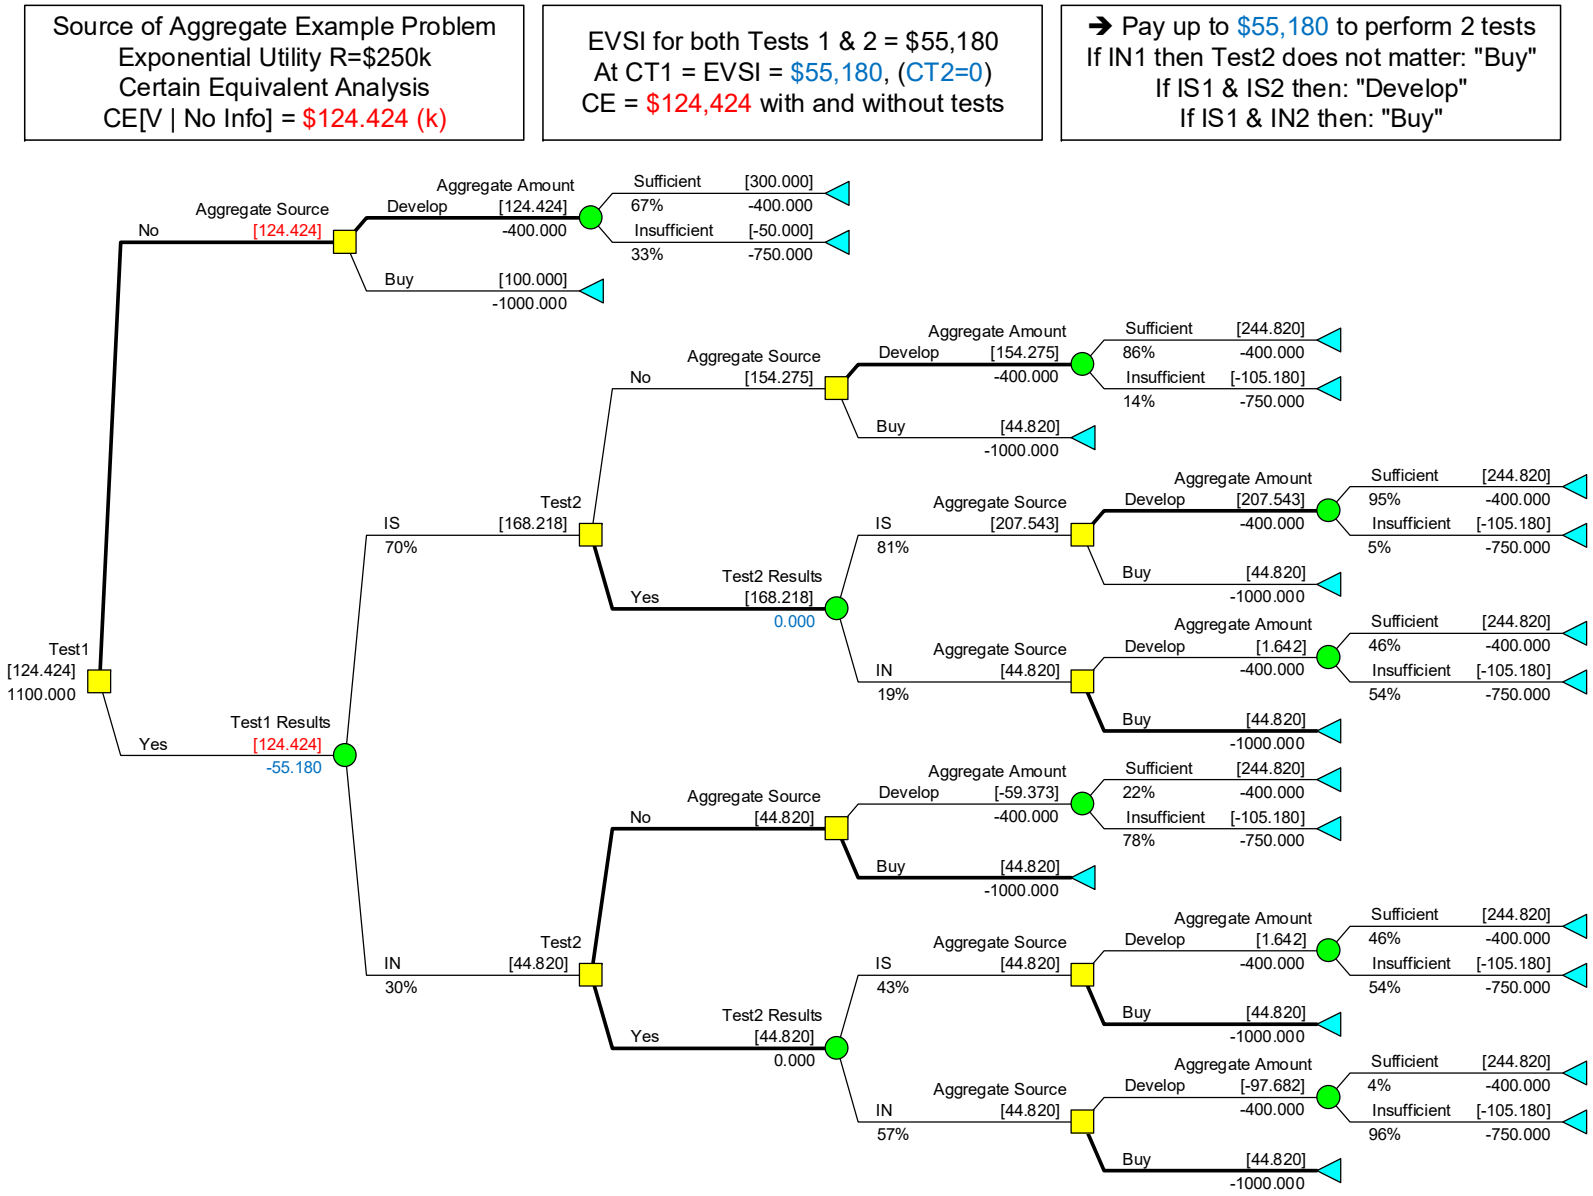


Figure U.34. CT1=EVSI for 2 Tests on Aggregate Amount using Exponential Utility, R=$250k

Source of Aggregate Example Exponential Utility R=$250k Certain Equivalent Analysis CE[V | No Info] = $124.424 (k)	If CT1=0 & IS1 ➔ EVSI2 for Test2=$13,493 If IN1 ➔ EVSI2 for Test 2 = $0 At CT2 = EVSI2 | IS1 = $13,943 EVSI1 = $171,318-124,424 = $46,894	EVSI1 (= maxCT1) depends on CT2: CT2=$0 ➔ maxCT1=$55,180 $0≤CT2≤$13,943 ➔ $46,894≤maxCT1≤$55,180 $13,943≤CT2 ➔ maxCT1=$46,894

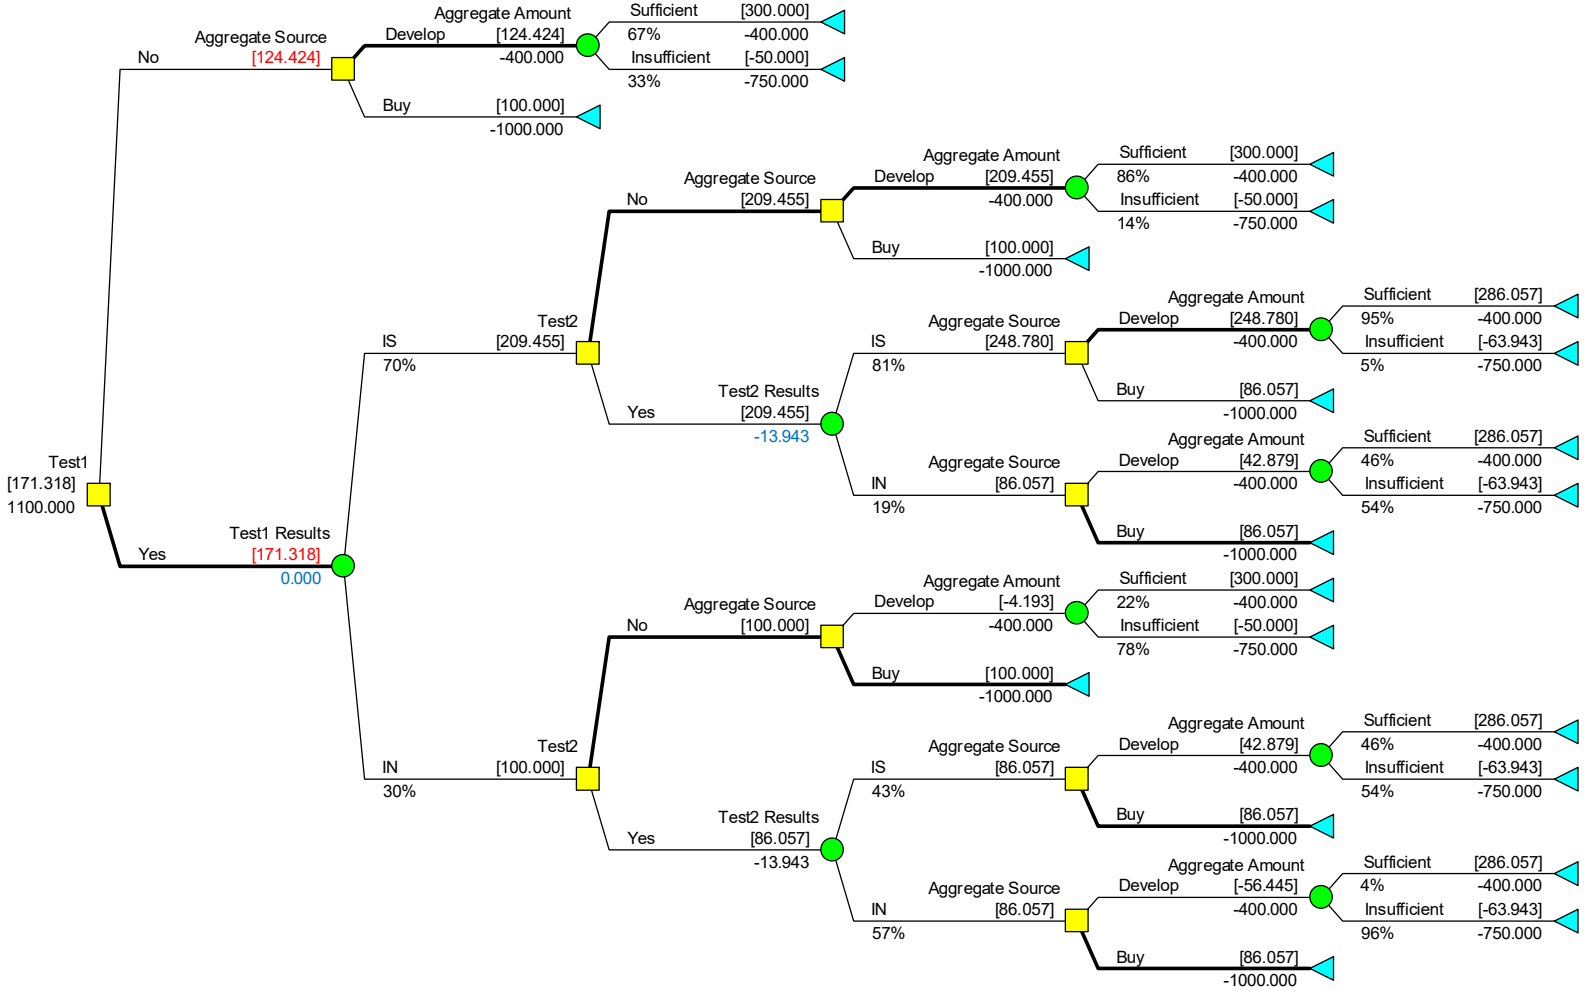


Figure U.35. EVSI for Test 1 & 2 on Aggregate Amount using Exponential Utility, R=$250k

Source of Aggregate Example Exponential Utility R=$250k Certain Equivalent Analysis CE[V | No Info] = $124.424 (k)	EVSI1 (= maxCT1) depends on CT2: CT2=$0 ➔ maxCT1=$55,180 $0≤CT2≤$13,943 ➔ $46,894≤maxCT1≤$55,180 $13,943≤CT2 ➔ maxCT1=$46,894	E.g., if CT2 = $13,943 ➔ EVSI1 =$46,894 At CT1 = EVSI1 = $46,894 & CT2 = $13,943 ➔ CE = $124,424 with and without tests

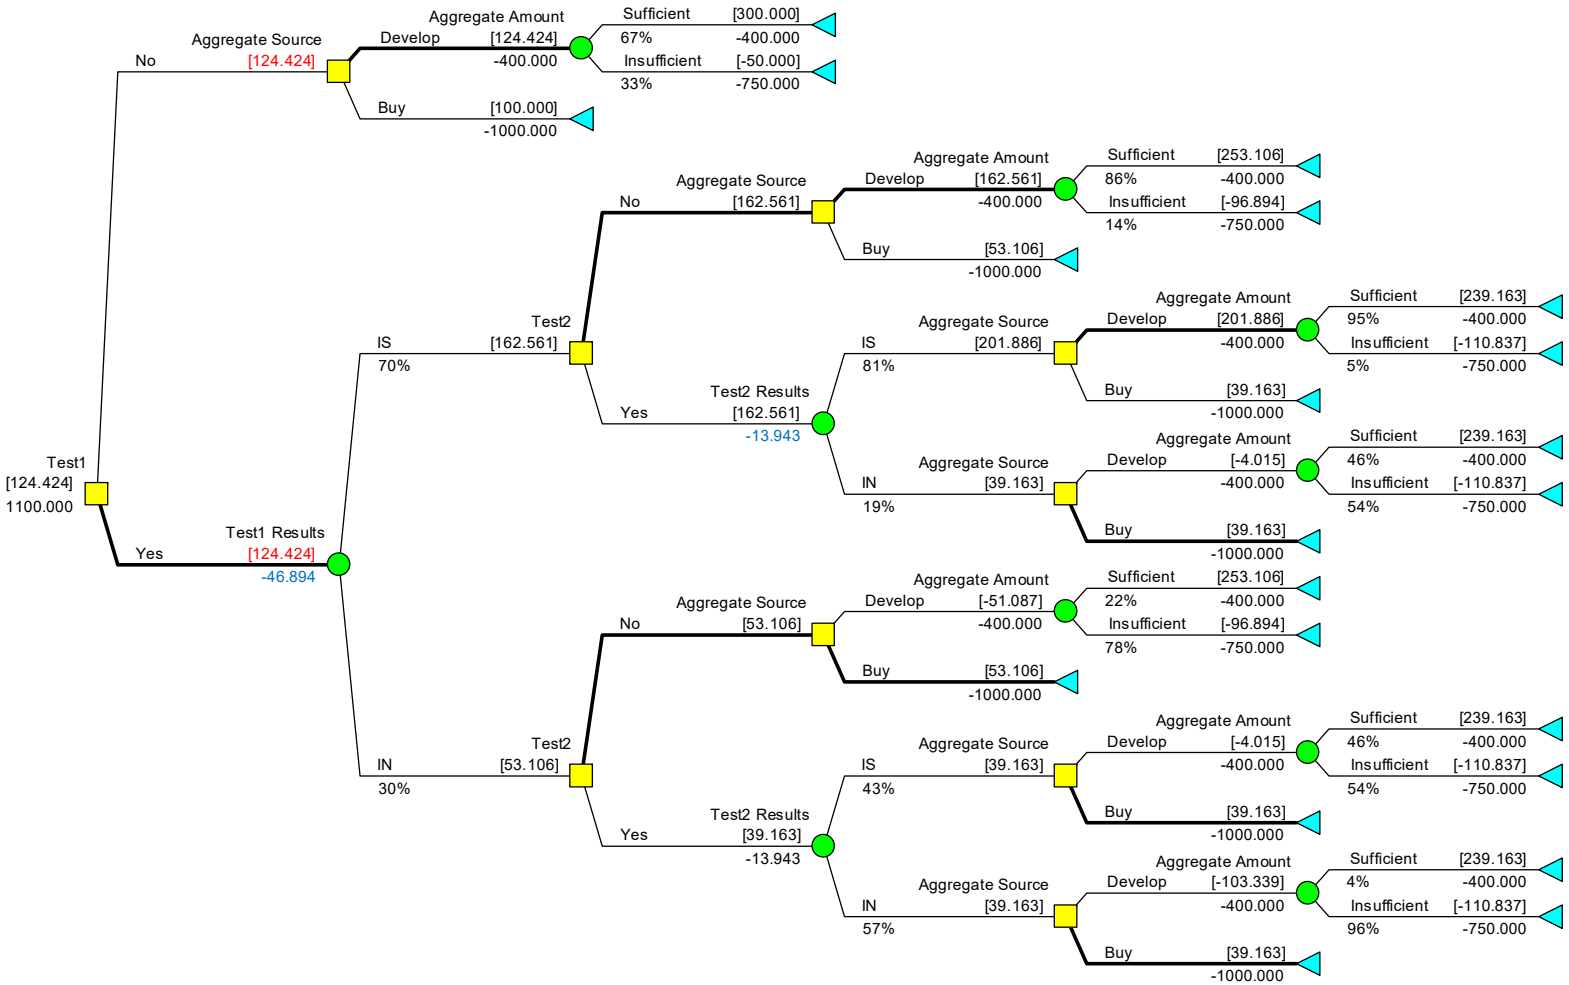


Figure U.36. EVSI for Test 1 | CT2 on Aggregate Amount using Exponential Utility, R=$250k

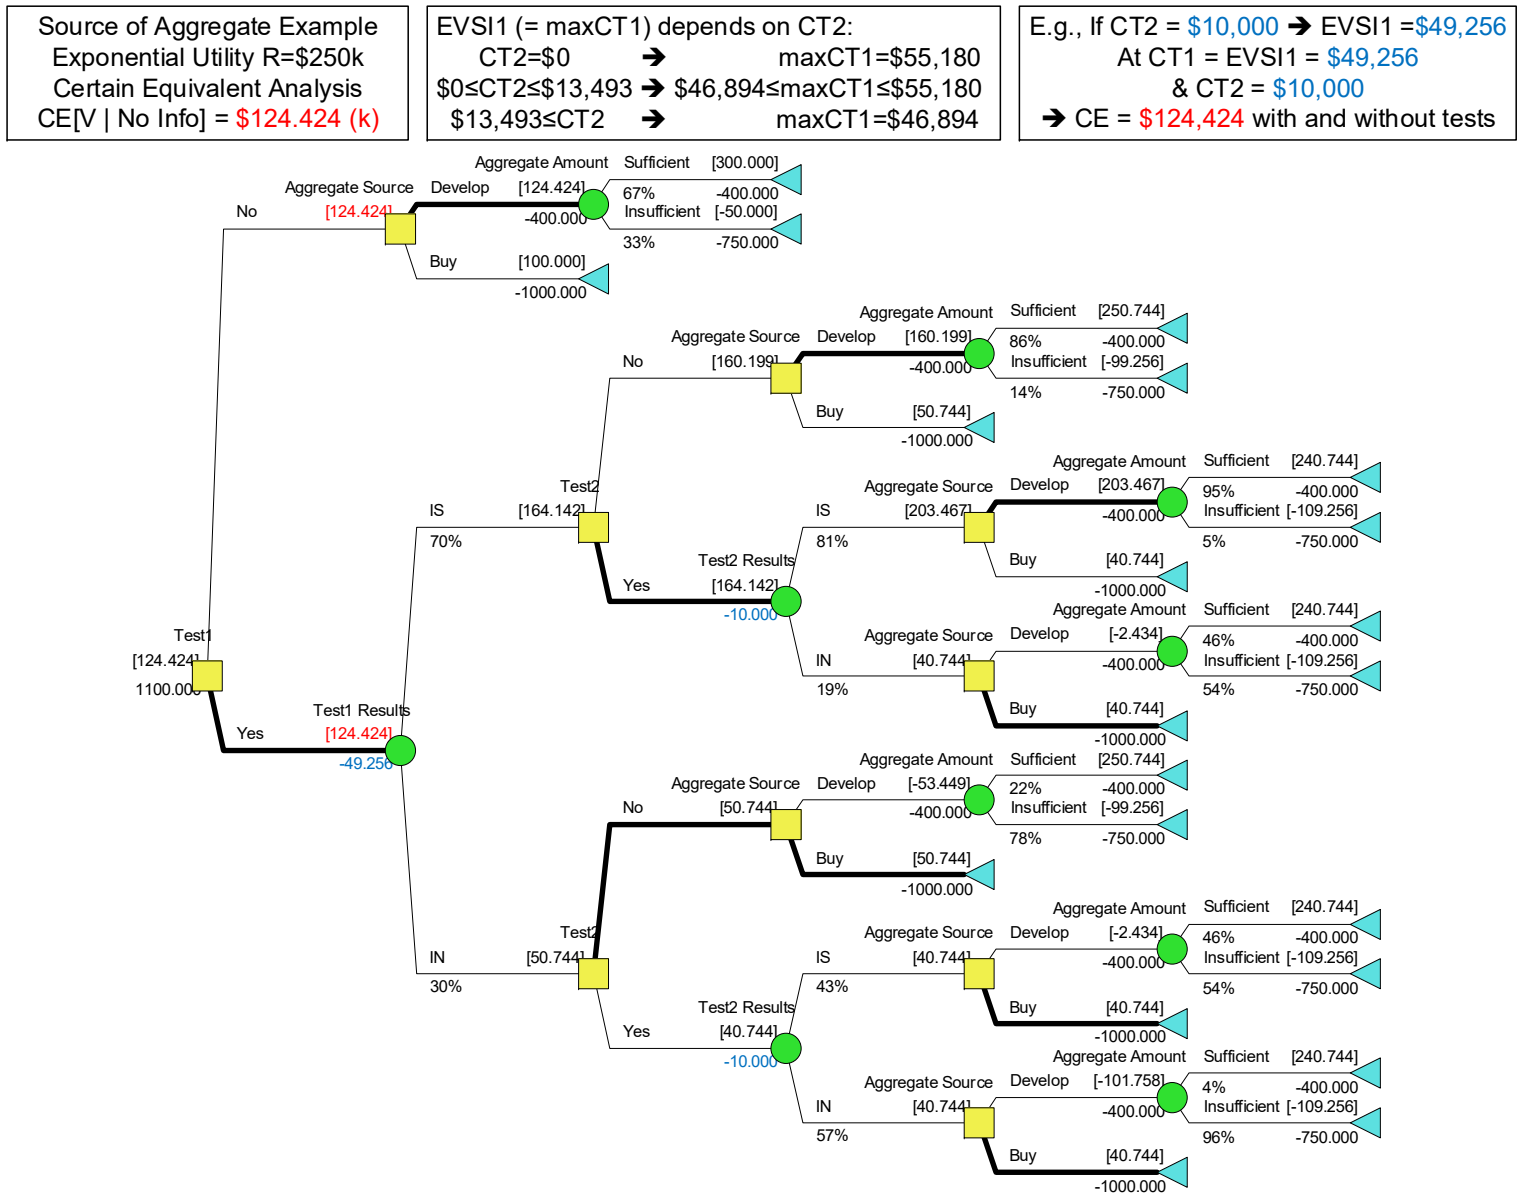
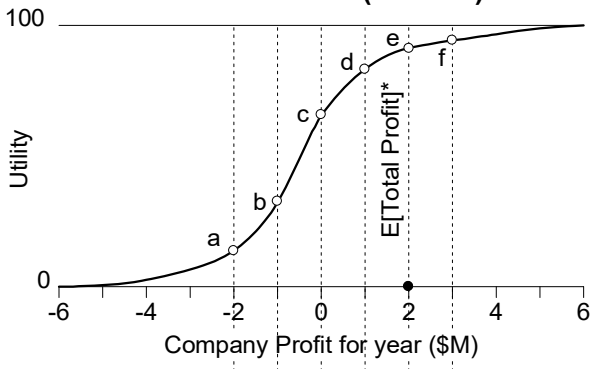
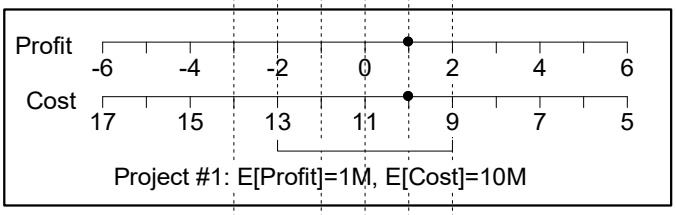


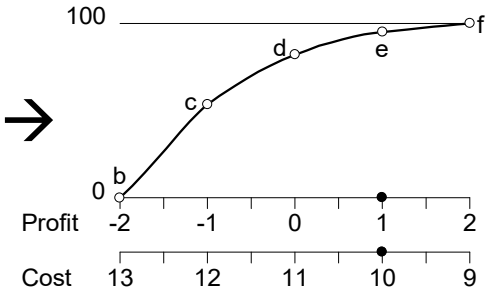
Figure U.36. EVSI for Test 1 | CT2 on Aggregate Amount using Exponential Utility, R=$250k

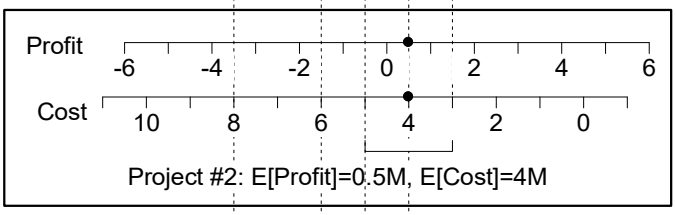
USING A COMPANY (PROJECT) UTILITY CURVE ON INDIVIDUAL PROJECTS (ACTIVITIES)

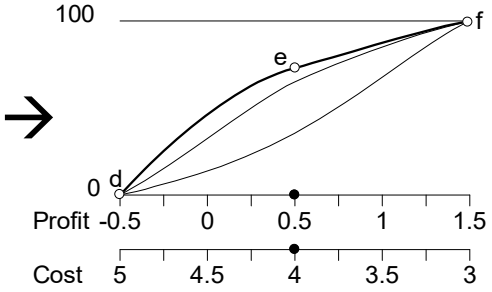


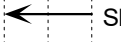
Utility curve is based on company's current asset position. E[Total Profit] is an estimate that includes profit from all projects for the year. The E[Profit] for any project is aligned with E[Total Profit] and the portion of the utility curve bounded by the range of the project's possible outcomes is used to make decisions for that project. This simple approximate method decouples projects so that they be analyzed individually. Notice that the utility curve indicates that the company may become a "gambler" when it nears bankruptcy.

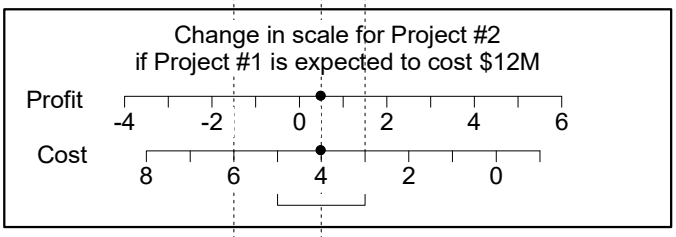


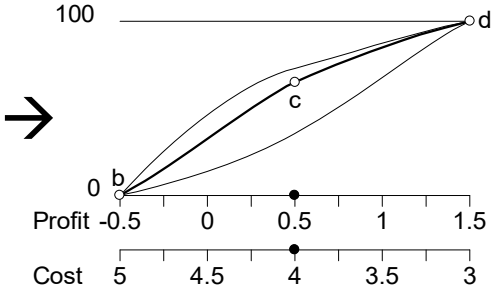


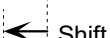


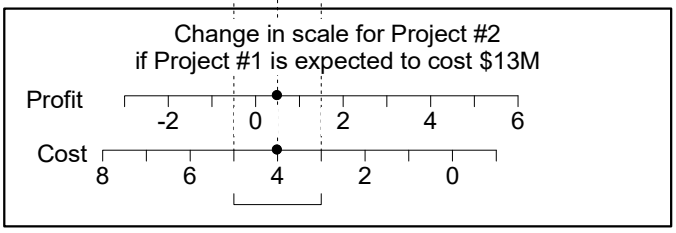


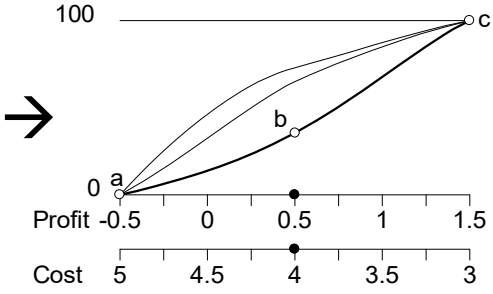












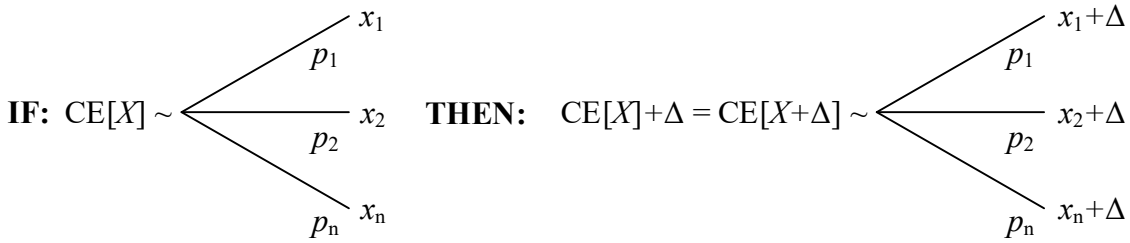
Exponential Utility Curve — Delta Property — Constant Risk Preference

Some decision makers are willing to accept what is known as the "*delta-property*", *i.e.*:

- "If all *prizes* of a lottery, X, are changed by the same constant amount *delta*, Δ, then the *certain equivalent* of the lottery is also changed by the same amount *delta*, Δ."

$$CE[X+\Delta] = CE[X] + \Delta$$

The *delta-property* is illustrated graphically by the following *indifference* statements:



Acceptance of the *delta-property* restricts the utility curve to be either *linear* or *exponential*:

Utility curve shape	**Function**	**Inverse function:**
Linear (risk neutral):	$u(x) = a + bx$	$x = [u(x) - a]/b$
Exponential:	$u(x) = a - be^{-\gamma x} = a - be^{-x/\rho}$	$x = -\rho \ln\{[a - u(x)]/b\}$

As explained earlier, the relative preferences expressed by a utility curve $u(x)$ do not change if we apply a *positive* linear transformation to obtain another utility curve $u'(x)$ as follows:

$$u'(x) = a' + b'u(x) \quad \text{where } b' > 0$$

Thus, the constants a and b for the exponential utility curve can be chosen arbitrarily so that its scale suits the decision problem being analyzed (similar to switching from a Fahrenheit scale to a Celsius scale). However, the constant b must have the appropriate sign so that utility increases in the preferred direction of x (the value measure).

The fact that a and b can be selected arbitrarily, leads to the important conclusion that the exponential utility function is completely defined by just one parameter, γ, or $\rho = 1/\gamma$:

γ = risk aversion coefficient		$\rho = 1/\gamma$ = risk tolerance	
Risk Averse Behavior:	$\gamma > 0$	$\rho > 0$	$b > 0$
Risk Preferring (*Gambler*):	$\gamma < 0$	$\rho < 0$	$b < 0$
Risk Neutral:	$\gamma \to 0$	$\rho \to \infty$	

A person who accepts the *delta*-property is sometimes called a "*delta-person*", a "*constant risk-averter*" or a "*person with constant risk-aversion*".

The *delta*-property is very useful in practice because it makes decision analysis calculations *significantly easier*. For this reason, we sometimes assume initially that the decision maker is a *delta*-person and has an exponential utility curve to make the preliminary investigation and sensitivity analysis much easier. The next sections and particularly the ease by which EVPI and EVSI can be calculated using the *delta*-property, should make this clear.

Properties of the Exponential Utility Curve

The exponential utility curve has the following useful properties:

- The "*Delta-Property*" (Δ = a constant; X = outcomes of a lottery = a random variable):

$$CE[X+\Delta] = CE[X] + \Delta$$

- The *PISP* (*personal indifferent selling price*) and the *PIBP* (*personal indifferent buying price*) of a lottery X are identical.

PISP:	$PISP = CE[X]$	Decision maker owns lottery X. By definition, $CE[X] = PISP$.
PIBP: (Δ property)	$CE[X - PIBP] = \$0$ $CE[X - PIBP + PIBP] = \$0 + PIBP$ $CE[X] = PIBP$	Decision maker does not own lottery X and is indifferent between paying PIBP to buy it and doing nothing ($0 change).
$\Rightarrow$	$PISP = PIBP$	Personal Indifferent Selling Price = Personal Indifferent Buying Price

- When a *delta*-person faces a sequence of *independent* lotteries, X_i, the certain equivalent for the whole prospect, X, is the sum of the certain equivalents for each individual lottery. (This is *not* true if X_1 and X_2 are *not* independent.)

$$X = X_1 + X_2 \qquad X_1, X_2 \text{ are independent lotteries}$$
$$CE[X] = CE[X_1] + CE[X_2]$$

- The certain equivalent, $CE[X]$, for a *profit* lottery, X, whose outcomes follow a Normal distribution with mean, m, and variance, σ^2, is given by the very useful formula:

$$CE[X] = m - \frac{1}{2}\gamma\sigma^2 = m - \frac{1}{2\rho}\sigma^2$$

- The risk premium, $RP[X]$, for a *profit* lottery, X, whose outcomes follow a Normal distribution with mean, m, and variance, σ^2, does not depend on m, but only on σ^2:

$$RP[X] = E[X] - CE[X] = m - m + \frac{1}{2}\gamma\sigma^2 = \frac{1}{2}\gamma\sigma^2 = \frac{1}{2\rho}\sigma^2$$

  By definition, for a risk-averse decision maker, the risk premium of *any profit* lottery, X, is *positive*. For a risk-averse *delta*-person, the risk-aversion coefficient, γ, and risk tolerance, ρ, are indeed positive. Thus, the risk premium of a *profit* lottery, X, is also positive:

$$\gamma > 0, \quad \rho > 0 \quad \Rightarrow \quad RP[X] > 0 \quad \Rightarrow \quad E[X] > CE[X]$$

  A *positive* risk premium, $RP[X]$, indicates that a risk-averse decision maker is willing to "forgo" (give up) a portion of the expected profit (i.e., the risk premium) and accept with certainty a smaller amount (i.e., the certain equivalent) to avoid the risk of the profit lottery.

- The risk premium, $RP[X]$, increases with risk aversion, i.e., as γ increases or ρ decreases. Also, the risk premium increases as the variance of the lottery increases (i.e., as the spread of the outcomes, or the "risk" of the lottery, increases).

Delta Property, Risk Premium and Constant Risk Aversion

A direct consequence of the *delta-property* is *constant risk-aversion*:

> ***Constant Risk Aversion***
> *The risk premium RP[X] of a lottery X remains the same and does not change even if all lottery prizes, X, are increased by the same amount Δ, to become X + Δ, or if the wealth of the decision-maker is increased or decreased by any amount Δ.*

The proof is quite simple:

$$
\begin{aligned}
E[X+\Delta] &= E[X]+\Delta && \text{: expected value of } X \text{ plus constant } \Delta \\
CE[X+\Delta] &= CE[X]+\Delta && \text{: } \textit{delta}\text{-property} \\
RP[X] &= E[X]-CE[X] && \text{: definition of } \textit{risk premium} \\
RP[X+\Delta] &= E[X+\Delta]-CE[X+\Delta] && \text{: definition of } \textit{risk premium} \\
&= E[X]+\Delta-CE[X]-\Delta && \text{: substitution} \\
&= E[X]-CE[X] && \text{: } \textit{deltas} \text{ cancel} \\
RP[X+\Delta] &= RP[X] && \text{: risk premium remains the same}
\end{aligned}
$$

The risk premium for a lottery is a measure of the degree of risk-aversion exhibited by a decision maker. The larger the risk premium for a given lottery, the more risk-averse the decision maker.

Other risk-averse utility curves typically exhibit *decreasing risk-aversion*. I.e., as the decision maker's wealth increases, his attitude towards a given lottery becomes less risk-averse and the risk premium decreases. For example, when the decision maker's total wealth is $100k, his risk premium for a particular lottery might be $1,000. If unexpectedly the decision maker inherits $100M, his risk premium for the same lottery might become much less than $1,000 (and might even approach zero) because the lottery outcomes and the risk involved in the lottery are not that important anymore.

This behavior towards risk is not true for the exponential utility curve. The degree of risk-aversion exhibited by a *delta*-person remains the same, irrespective of changes in that person's wealth. Put differently, a delta-person exhibits *constant risk-aversion*.

Constant risk-aversion is a direct consequence of the *delta*-property. Because of the acceptance of the *delta*-property, the risk premium for a lottery does not change when all outcomes are increased by the same amount Δ (an increase in wealth), even if Δ is very large, such as $100M. Hence, the degree of risk-aversion exhibited by a *delta*-person with an exponential utility curve is *constant* and *independent of the decision maker's wealth*. And that is one of the reasons why a person with an exponential utility curve is sometimes referred to as a "*constant risk-averter*" in the literature.

Using the Delta Property to compute EVPI and EVSI directly.

The concepts below apply equally well to (a) *perfect* information and the calculation of *EVPI* (Expected Value of Perfect Information) and (b) *sampled* information and *EVSI* (Expected Value of Sampled Information). Thus, the shorthand notation {perfect | sampled} is used to refer to either perfect or sampled information and {EVPI | EVSI} to refer to either EVPI or EVSI.

Consider a general decision problem where the *maximum* expected utility among all decision alternatives, *excluding* the alternative to obtain {perfect | sampled} information, is given by $E[u(X) \mid \text{no info}]$. Here the end-node outcomes of the decision tree are represented by X.

By definition, the corresponding *certain equivalent*, $CE[X \mid \text{no info}]$, for the best decision alternative, is that amount of money that has the same utility as $E[u(X) \mid \text{no info}]$:

$$u(CE[X \mid \text{no info}]) = E[u(X) \mid \text{no info}]$$

Suppose that {perfect | sampled} information can be acquired for *free*, i.e., CT=\$0, where CT = Cost of the Test = cost of obtaining the information (determined by the market).

In general, the decision branch to acquire {perfect | sampled} information leads to several end-nodes. Let X_i be the dollar values of these end-node outcomes when CT=\$0.

When the decision tree is analyzed, the decision branch to obtain {perfect | sampled} information has an expected utility of $E[u(X) \mid \text{info}, CT=0]$ under the assumption that CT=\$0.

The amount of money that has the same utility as $E[u(X) \mid \text{info}, CT=0]$ is, by definition, the corresponding *certain equivalent*, $CE[X \mid \text{info}, CT=0]$.

$$u(CE[X \mid \text{info}, CT=0]) = E[u(X) \mid \text{info}, CT=0]$$

In general, the option to acquire information at zero cost, and which does not eliminate any of the otherwise available decision alternatives, cannot have negative value. This means that:

$$E[u(X) \mid \text{info}, CT=0] \geq E[u(X) \mid \text{no info}]$$

$$CE[X \mid \text{info}, CT=0] \geq CE[X \mid \text{no info}]$$

The {EVPI | EVSI} of {perfect | sampled} information is defined to be that cost of acquiring the information, CT^*, that makes the decision maker *indifferent* between obtaining and not obtaining the information. Mathematically:

$$E[u(X\text{-}CT^*) \mid \text{info}, CT^*] = E[u(X) \mid \text{no info}]$$

Or, equivalently,

$$CE[X\text{-}CT^* \mid \text{info}, CT^*] = CE[X \mid \text{no info}]$$

I.e., the {EVPI | EVSI} of {perfect | sampled} information equals that special cost for acquiring the information, CT^*, that results in the same {expected utility | certain equivalent} "*with*" and "*without* the {perfect | sampled} information". (Of course, there is a different CT* for the EVPI and a different CT* for the EVSI because EVPI ≥ EVSI.) Discussed below is the equality of the certain equivalents "*with*" and "*without* information, when the *delta property* applies.

Implications of the Delta Property

When the cost of acquiring {perfect | sampled} information is *CT**=*PIBP*, all end-node dollar values X_i for the branch associated with obtaining {perfect | sampled} information are reduced by the same amount *delta*, $\Delta = CT^*$, and become X_i-*CT**.

The *delta*-property of the exponential utility curve asserts that:

If all prizes of a lottery are changed by *delta*,
then the certain equivalent of the lottery is also changed by *delta*.

In this case, the *delta property* means that,

If all the end-node dollar outcomes X_i are reduced by *CT**
then the certain equivalent (that corresponds to the prizes X_i) is also reduced by *CT**.

The following indifference statement is true because of the definition of *PIBP*= *CT**:

$$CE[X\text{-}CT^* \mid \text{info}, CT^*] = CE[X \mid \text{no info}]$$

And by the *delta property*:

$$CE[X\text{-}CT^* \mid \text{info}, CT^*] = CE[X \mid \text{info}, CT=0] - CT^*$$

Because the left-hand sides are equal, the right-hand sides must also be equal:

$$CE[X \mid \text{info}, CT=0] - CT^* = CE[X \mid \text{no info}]$$

Thus, the value of information CT* is given by:

$$CT^* = CE[X \mid \text{info}, CT=0] - CE[X \mid \text{no info}]$$

From here we conclude that:

$CT^* = CE[X \mid \text{info}, CT=0] - CE[X \mid \text{no info}]$
$EVPI = CE[X \mid \text{perfect info}, CT=0] - CE[X \mid \text{no info}]$
$EVSI = CE[X \mid \text{sampled info}, CT=0] - CE[X \mid \text{no info}]$

Conclusion: For a delta-person, the {EVPI | EVSI} is the monetary difference between the Certain Equivalents (*CE*) "*with*" and "*without*" ***free*** {perfect | sampled} information.

The two certain equivalents on the right-hand side "*with*" and "*without*" ***free*** information can be computed directly and easily. Their difference provides *CT** which equals EVPI or EVSI depending on whether the information to be obtained is perfect or sampled.

The value of the above results in practice is significant. For all other risk-sensitive utility curves the determination of {EVPI | EVSI} requires repeated trial-and-error calculations that, unless performed by computer, can be quite laborious. As shown above, because of the delta-property, the calculation of {EVPI | EVSI} for an exponential utility curve does not require trial-and-error and is almost as simple and straightforward as that for the risk-neutral case.

© 2024 *Photios G. Ioannou*

General Approximations Based on Properties of Exponential Utility Curves

As stated above, when the utility curve is exponential, the certain equivalent for a *profit* lottery, X, whose outcomes follow a Normal distribution with mean, m, and variance, σ^2, is:

$$CE[X] = m - \frac{1}{2}\gamma\sigma^2 = m - \frac{1}{2\rho}\sigma^2$$

And the risk premium for this lottery is given by:

$$RP[X] = \frac{1}{2}\gamma\sigma^2 = \frac{1}{2\rho}\sigma^2$$

Empirical studies have shown that the above formulas for the certain equivalent and risk premium are *good approximations for all lotteries*, including cases where:

- The lottery is continuous, but follows a distribution other than the Normal
- The lottery is discrete and even binary.

When does Risk Aversion matter on Ranking Decision Alternatives?

Given its wide applicability as an approximation, the above formula for the risk premium can also be used to identify *the types of decisions where risk aversion matters*: I.e., where risk-aversion changes the optimal decision alternative from that found by expected monetary value.

Let us start with an important empirical observation made by Howard[1],

> *"While the ability to capture risk preference is an important part of our conceptual view of decision-making, I find it is a matter of real practical concern in only 5 percent to 10 percent of business decision analysis. Of course, the situations that require risk preference, such as bidding or portfolio problems, use it very seriously."*

At least a partial explanation of Howard's findings is provided by the above approximation for the risk premium. Since the risk premium is the difference between the expected value and the certain equivalent of a lottery:

> *If the various decision alternatives from which to choose are associated with lotteries that have comparable risk premiums, then the ranking of decision alternatives based on expected monetary value is the same as their ranking based on certain equivalents.*
>
> *Risk attitude will impact the ranking of decision alternatives only if the risk premiums are significantly different from one decision alternative to the next.*

Hence, risk-aversion matters only when the approximate risk premiums, $\gamma\sigma^2/2$, are significantly different from one decision alternative to the next. Otherwise, the term $\gamma\sigma^2/2$ is like a constant correction to the expected value of each decision alternative and ranking based on expected values is the same as ranking based on certain equivalents.

Clearly, the impact of risk attitude increases as the decision maker becomes *more risk-averse* and as the *relative difference of the variance of value* (*profit*) among decision alternatives *increases*. The combination of these two effects increases the relative difference between risk premiums and consequently their impact on ranking decision alternatives.

[1] Howard, R. A. 1988. Decision Analysis: Practice and Promise. Management Science, 34, 679-695.

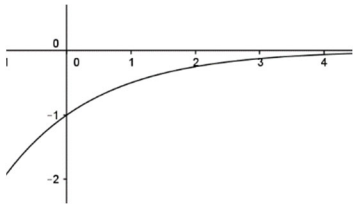
Functional Forms for the Risk-Averse Exponential Utility Curve — Scaling

To simplify the analysis of a particular problem, we often make assumptions about the functional form for the risk-averse exponential utility curve to change it from its standard form:

$$u(x) = a - be^{-\gamma x}$$

Constants *a* = 0 and *b* =1

The most basic form sets $a = 0$ and $b = 1$ and is often used for theoretical analysis:



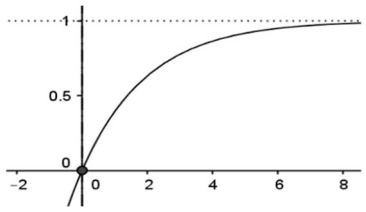
$$u(x) = -e^{-\gamma x} = -e^{-x/\rho}$$

$$x = -\rho \ln[-u(x)]$$

This form is not used often in practice because the utility function, $u(x) < 0$, is *negative* for all values *x*. It approaches $u(x) = 0$ asymptotically as $x \to \infty$.

Constants *a* and *b* equal to 1

A classic form is to set both constants *a* and *b* equal to 1:



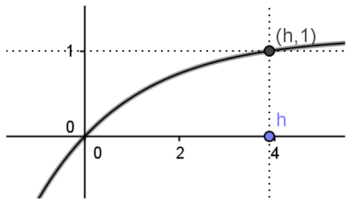
$$u(x) = 1 - e^{-\gamma x} = 1 - e^{-x/\rho}$$

$$x = -\frac{1}{\gamma}\ln[1-u(x)] = -\rho \ln[1-u(x)]$$

This form has the advantage that $u(0) = 0$ and approaches $u(x) = 1$ asymptotically as $x \to \infty$.

Scaled from $0 to $h

Often it is desirable for the exponential utility function to be scaled between a lower bound $u(0) = 0$ and an upper bound $u(h) = 1$. In this case, it takes the functional form:

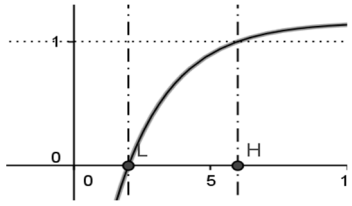


$$u(x) = \frac{1-\exp(-\gamma x)}{1-\exp(-\gamma h)} = \frac{1-\exp(-x/\rho)}{1-\exp(-h/\rho)}$$

$$x = -\rho \ln\{1 - u(x)[1 - \exp(-h/\rho)]\}$$

Scaled from $L to $H

In the more general case where the lower bound is $u(L) = 0$ and the upper bound is $u(H) = 1$ the exponential utility function is written as:



$$u(x) = \frac{1-\exp(-\gamma(x-L))}{1-\exp(-\gamma(H-L))} = \frac{1-\exp(-(x-L)/\rho)}{1-\exp(-(H-L)/\rho)}$$

$$x = -\rho \ln\{1 - u(x)[1 - \exp(-(H-L)/\rho)]\} + L$$

Certain Equivalent of Exponential Utility is Independent of Functional Form

For a *delta*-person with an exponential utility curve, the certain equivalent $\tilde{x} = CE[X]$ for lottery X does not depend on which functional form for the exponential utility is actually used. This means that the certain equivalent $\tilde{x} = CE[X]$ is always given by the *same formula* (shown in the box below) for any of the functional forms shown above. This is an important result.

Using the most general functional form for the exponential utility curve, the formula for the certain equivalent $\tilde{x}$ for a lottery X can be derived as follows:

$$u(\tilde{x}) = E[u(x)]$$

$$\frac{1-\exp(-(\tilde{x}-L)/\rho)}{1-\exp(-(H-L)/\rho)} = E\left[\frac{1-\exp(-(x-L)/\rho)}{1-\exp(-(H-L)/\rho)}\right]$$

$$1-\exp(-(\tilde{x}-L)/\rho) = E[1-\exp(-(x-L)/\rho)]$$

$$\exp(-(\tilde{x}-L)/\rho) = E[\exp(-(x-L)/\rho)]$$

$$\exp(L/\rho)\exp(-\tilde{x}/\rho) = \exp(L/\rho)E[\exp(-x/\rho)]$$

$$\boxed{\tilde{x} = CE[X] = -\rho \ln E[\exp(-x/\rho)]}$$

Thus, to find the certain equivalent $\tilde{x}$ for a lottery X we need to find the *expected value* of the exponential *function* of x, i.e., $\exp(-x/\rho)$, take the natural logarithm of the result and multiply it by the negative risk tolerance. None of the constants a, b, L, H, are part of the formula.

A very important application of the above certain equivalent of the exponential utility function $\tilde{x} = CE[X]$ is when X follows a Normal distribution with mean m and variance σ^2:

$$\boxed{\tilde{x} = CE[X] = m - \frac{1}{2}\gamma\sigma^2 = m - \frac{1}{2\rho}\sigma^2}$$

Risk Tolerance — Appropriate Order of Magnitude

In practice, it is often preferred to express the function for the exponential utility curve using risk tolerance, ρ, as opposed to the risk aversion coefficient, γ. The reason is that risk tolerance has the same units as the attribute, x, (e.g., dollars). As a result, it is much easier to understand the meaning of risk tolerance, ρ, and the implications of its magnitude.

> For a meaningful analysis, the risk tolerance, ρ, should be of the same order of magnitude as the interval width from the lower bound L to the upper bound H.

- When risk tolerance is very small, $\rho < 0.1(H - L)$, then the implied behavior becomes very risk-averse, and the assumption of *constant risk aversion* becomes questionable.
- When risk tolerance is large, $\rho > 10(H - L)$, the utility function is approximately linear, and decisions could be based on expected monetary value (EMV) alone.

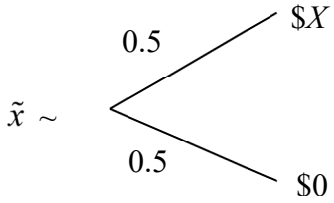
The question of the proper magnitude for risk tolerance, ρ, is very important and is addressed in more detail below when facing large monetary prospects.

Risk Tolerance and Large Monetary Prospects

Consider a person with an exponential utility function, scaled so that $u(\$0) = 0, \quad u(\$100) = 1$:

$$u(x) = \frac{1 - \exp(-x/\rho)}{1 - \exp(-100/\rho)}, \text{ where } \rho = \$72.13$$

Suppose that this person faces the following lottery where $\$X$ is the desirable (high) prize:

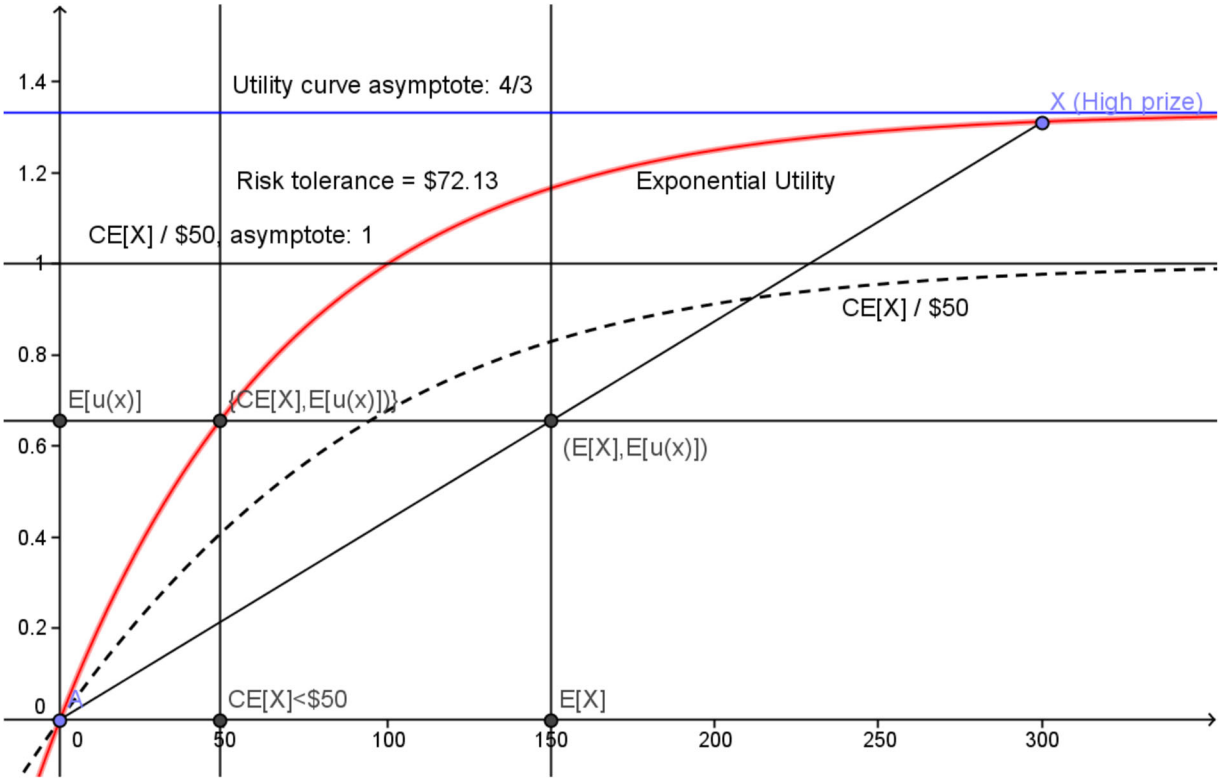


The utility of the certain equivalent $\tilde{x}$ equals the expected utility of the lottery:

$$u(\tilde{x}) = 0.5u(0) + 0.5u(X) = 0.5\frac{1 - \exp(-X/\rho)}{1 - \exp(-100/\rho)}$$

An example graph of this utility curve is shown in the figure below where:

- The lottery prizes are \$0 and X=\$300
- The expected monetary value is E[X]=\$150
- The corresponding certain equivalent is $\tilde{x}$ =CE[X] = \$48.88.



We can find the certain equivalent $\tilde{x}$ as a function of the prize \$$X$ by inverting the utility curve:

$$\tilde{x} = u^{-1}\left(0.5u(0) + 0.5u(X)\right) = u^{-1}\left(0.5u(X)\right)$$

$$u(\tilde{x}) = \frac{1-\exp(-\tilde{x}/\rho)}{1-\exp(-100/\rho)} = 0.5\frac{1-\exp(-X/\rho)}{1-\exp(-100/\rho)}$$

$$\tilde{x} = [\ln 2 - \ln(1+\exp(-X/\rho))]\rho$$

As shown by the dashed line in the graph, the function $\tilde{x}/50$ has an upper limit represented by the asymptote y=1.

This means that the certain equivalent $\tilde{x}$ for this 50-50 lottery will never exceed \$50 irrespective of how large the high prize \$$X$ becomes.
Thus, this decision maker would be willing to sell a 50-50 lottery with prizes (\$0, \$300) for a mere \$50.
And he would also be willing to sell the lottery (\$0, \$100Million) for the same \$50.
Very few people (if any) would accept this behavior as reasonable.

Notice that this particular exponential utility curve has a horizontal asymptote at 4/3 and that it reaches approximately 98% of its maximum value after only four times the risk tolerance (\$288). As a result, the utility, $u(X)$, is very close to 4/3 for any value, X, greater than about \$300. From the curve, we also see that the utility of zero is equal to zero. Therefore, the expected utility value of this lottery for any large outcome X > \$300 is very close to (1/2)*(0) + (1/2)*(4/3) = 2/3. Thus, for any large outcome X > \$300 the certain equivalent is about $\tilde{x} \approx u^{-1}(2/3) \approx \$50$.

The main reason for this phenomenon is that for large values of X the exponential utility curve becomes virtually constant and equals about 4/3. As a result, the expected utility of the lottery (and hence the utility of the certain equivalent) is close to 2/3 for any large value of X > \$300.

Risk Tolerance and the Certain Equivalent Limit for a 50-50 Lottery

It is not hard to find the upper limit of a delta-person's certain equivalent $\tilde{x}$ (= PISP = PIBP) for the above lottery as a function of risk tolerance ρ. For this purpose, it is convenient to write the exponential utility function in its simplest possible form:

$$u(x) = -\exp(-x/\rho)$$

For a large X, the expected utility of the lottery (i.e., the utility of the certain equivalent) is:

$$u(\tilde{x}) = -\exp(-\tilde{x}/\rho) = 0.5u(0) + 0.5u(X) \approx 0.5(-1) + 0.5(0) = -0.5$$

$$\exp(-\tilde{x}/\rho) = 0.5$$

$$\boxed{\tilde{x} = \rho \ln 2 = 0.693\rho}$$

Thus, the certain equivalent $\tilde{x}$ will never exceed about 70% of ρ, even for a very large high prize, X. In the example above, we have ρ = \$72.13 which implies that $\tilde{x} < \$50 \approx 0.693 \times 72.13$.

Risk Tolerance and the Certain Equivalent Limit for a General Lottery

Let us now consider a general simple lottery ($*X*, *p*; $0) where the probability of winning the high prize *X* is *p* and the probability of winning zero is 1-*p*. For a large *X*, the expected utility of the lottery (i.e., the utility of the certain equivalent) is:

$$u(\tilde{x}) = -\exp(-\tilde{x}/\rho) = (1-p)u(0) + pu(X) \approx (1-p)(-1) + p(0) = -(1-p)$$

$$\exp(-\tilde{x}/\rho) = (1-p)$$

$$\boxed{\tilde{x} = -\rho \ln(1-p)}$$

Thus, the certain equivalent $\tilde{x}$ will never exceed $-\rho \ln(1-p)$ no matter how large *X* is. The table below shows the factor $-\ln(1-p)$ for the risk tolerance ρ, for values of the probability *p*:

p	$-\ln(1-p)$	$\tilde{x}$	*E*[*X*]
0.5	0.693	$< 0.7\,\rho$	0.5*X*
0.1	0.105	$< 0.1\,\rho$	0.1*X*
0.01	0.01005	$< 0.01\,\rho$	0.01*X*
0.001	0.001001	$< 0.001\,\rho$	0.001*X*
0.0001	0.0001	$< 0.0001\,\rho$	0.0001*X*

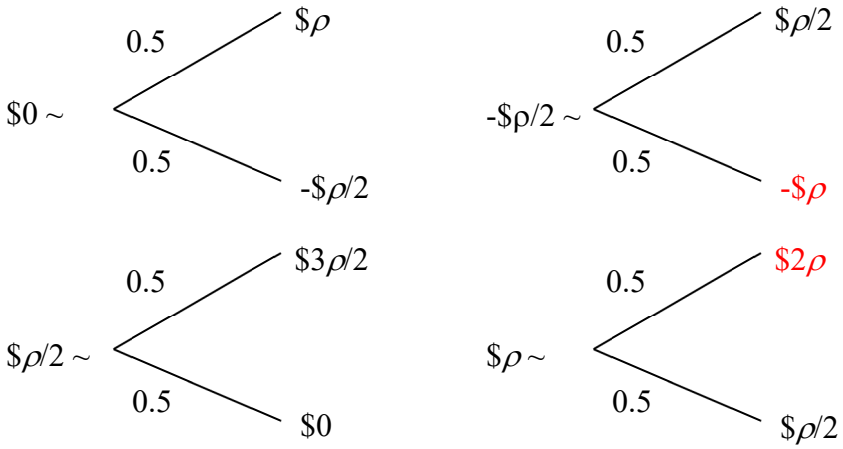
Thus, for probability values $p < 10\%$, the certain equivalent of the lottery cannot exceed $p\rho$, even though the expected monetary value of the lottery, pX, can be arbitrarily large. So, no matter how large $*X* is, the certain equivalent for the lottery ($*X*, *p*; $0) has an upper bound of:

$$\boxed{\tilde{x} <\approx p\rho \quad \text{for } p < 10\%}$$

Risk Tolerance and the Range of an Exponential Utility Curve

From this discussion, it is evident that the *range of monetary values* for an exponential utility curve should be selected with care if it is to serve as a reasonable model for human behavior.

As shown in a section below, we can assess a *delta*-person's approximate risk tolerance ρ by finding the payoffs $-\rho/2$ and ρ that would make the person indifferent between (a) receiving nothing and (b) a 50-50 lottery with payoffs of $-\rho/2$ and ρ.



The application of the *delta* property above, indicates that this person would have a certain equivalent of $-\rho/2$ for a 50-50 lottery with payoffs $-\rho$ and $\rho/2$. Furthermore, the person would have a certain equivalent of $\rho/2$ for a 50-50 lottery with payoffs 0 and $3\rho/2$. Finally, the person would have a certain equivalent of ρ for a 50-50 lottery with payoffs $\rho/2$ and 2ρ.

These results seem quite reasonable. As a result, a prudent *delta*-person should be careful when applying his utility curve to lotteries with payoffs that lie outside the range $-\rho$ to 2ρ, or to decisions where the range of possible outcomes exceeds 3ρ.

Assessing the Exponential Utility Function using Risk Odds

An easy way to encode a person's exponential utility function is through a simple lottery and the assessment of *risk odds* as described below. To establish a relationship between the exponential utility function and *risk odds* we use the following functional form:

$$\textit{Risk Averse}: \quad u(x) = a - br^{-x} \qquad r > 1,\ b > 0,\ -b < 0$$

$$\textit{Risk Preferring}\ (\textit{Gambler}): \quad u(x) = a + br^{-x} \qquad r < 1,\ b > 0$$

$$\textit{Risk Neutral}: \quad u(x) = a + bx \qquad r = 1,\ b > 0$$

The constant a can have any sign. For u (utility) to be increasing with x (e.g., money), b must have the sign indicated above. The traditional functional form is written with a minus sign before the coefficient b to accommodate the more common case of a risk-averse person dealing with money x. In this case, a positive number b automatically makes $-b$ negative, which is what it should be for risk-averse behavior.

Risk Odds

Consider a person with an exponential utility function

$$u(x) = a - br^{-x}$$

and an indifference statement where \$0 is the certain equivalent of a lottery with prizes \$1, -\$1:

$$\$0 \sim (\$1,\ p;\ -\$1,\ 1\text{-}p)$$

For *any* person, there is a certain probability value p that makes this indifference statement true.

It is not difficult to show that for a person with an *exponential* utility curve (a *delta* person), the *odds* that establish the above indifference statement, called *risk odds*, equal the coefficient r of the exponential utility curve:

$$\textit{Risk Odds} = r = \frac{p}{1-p}$$

Thus, we can encode the entire utility function for a *delta* person by using this simple lottery and a probability wheel to establish the probability p that makes the person indifferent between the lottery and \$0. Knowing the probability p gives the *risk odds*, which equal the parameter r that defines the exponential utility function completely. The constants a and b can be chosen arbitrarily to scale the utility function as needed.

The above equation gives the *risk odds* r if we know the probability p. It is easy to show that if we know the risk odds r, we can calculate the probability p as:

$$p = \frac{r}{1+r}$$

Generalized Risk Odds

In some cases, the prizes of the above lottery, \$1 and -\$1, are too small and too close to each other to encode the probability *p* accurately. To address this problem, we can use a similar lottery where the prizes have been scaled by a factor *m*.

Let us now generalize the lottery while keeping the certain equivalent at \$0. Consider an equivalence statement for a *delta* person where the lottery prizes are \$*m* and -\$*m*:

$$\$0 \sim (\$m,\ q;\ -\$m,\ 1\text{-}q)$$

The probability *q* is such that \$0 is the certain equivalent of the lottery with prizes \$*m*, -\$*m*.

The *risk odds* for this lottery that make the above indifference statement true are written as *r*(*m*) to emphasize that the lottery prizes are \$*m* and -\$*m.*

$$Risk\ Odds = r(m) = \frac{q}{1-q}$$

It can be shown easily that the *risk odds r*(*m*) for this lottery are also related to the parameter *r* of the exponential utility function:

$$Risk\ Odds = r(m) = r^m = \frac{q}{1-q} = \left(\frac{p}{1-p}\right)^m$$

We can restate the above equation for $r = r(1)$ as:

$$r = \frac{p}{1-p} = r(m)^{\frac{1}{m}} = \left(\frac{q}{1-q}\right)^{\frac{1}{m}}$$

The parameter *r* of the exponential utility function is the m^{th} root of the risk odds *r*(*m*).

This equation also allows us to relate the risk odds *r*(*m*) of a lottery with prizes $\pm\$m$ and probability *q* to the risk odds *r*(*n*) of a lottery with prizes $\pm\$n$ and probability *s*:

$$r = \frac{p}{1-p} = r(m)^{\frac{1}{m}} = \left(\frac{q}{1-q}\right)^{\frac{1}{m}} = r(n)^{\frac{1}{n}} = \left(\frac{s}{1-s}\right)^{\frac{1}{n}}$$

Functional Form Parameters and Risk Preference

	Risk Averse	Risk Neutral	Risk Preferring
Risk Odds, r	$r > 1$	$r = 0$	$r < 1$
Risk Aversion, γ	$\gamma > 0$	$\gamma = 0$	$\gamma < 0$
Risk Tolerance, ρ	$\rho > 0$	$\rho = \infty$	$\rho < 0$

Encoding an Exponential Utility Function using Risk Odds

Suppose that a *delta*-person has stated risk odds $r(10)=3$ for a lottery where one wins \$10 with probability q =75% or loses \$10 with probability 1-q=25%. This means that:

$$q = 75\% = \frac{r}{1+r} = \frac{3}{1+3}$$

$$r(10) = \frac{q}{1-q} = 3 = r^{10} = (\frac{p}{1-p})^{10}$$

$$r = \sqrt[10]{3} = 1.116123$$

$$p = \frac{1.116123}{1+1.116123} = 0.527438 = 53\%$$

Notice that it is much easier to assess q=75% accurately than to assess p=53% accurately.

These results imply that this person's exponential utility curve is:

$$u(x) = a - br^{-x} = a - b(1.116123)^{-x}$$

Exponential Utility Functions: Standard Form and Risk Odds Form

As shown above, the risk odds form of the exponential utility curve with parameter r is:

$$u(x) = a - br^{-x}$$

Whereas the standard form with parameter γ is written as:

$$u(x) = a - be^{-\gamma x}$$

As mentioned earlier, γ = risk aversion coefficient and its reciprocal $\rho = \frac{1}{\gamma}$ = risk tolerance.

From here we see that

$$\gamma = \ln(r)$$

$$\rho = \frac{1}{\gamma} = \frac{1}{\ln(r)}$$

Continuing the example above, where

$$r = 1.116123$$

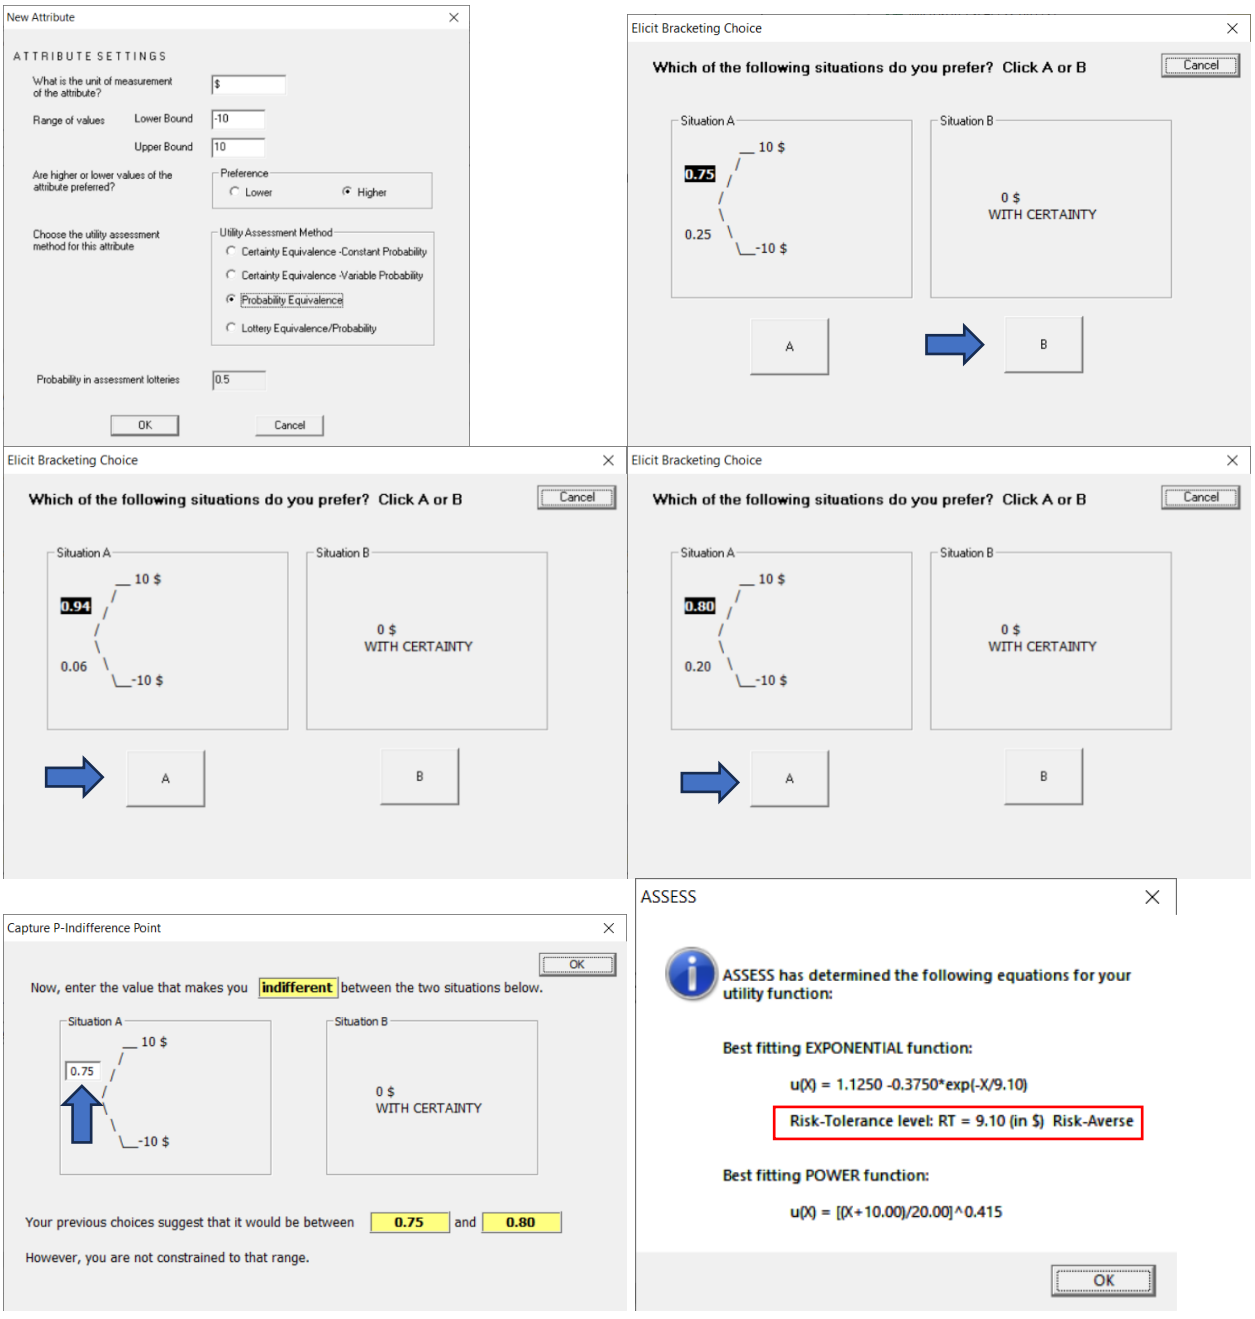
$$\gamma = \ln(r) = \ln(1.116123) = 0.109861/\$$$

$$\rho = \frac{1}{\gamma} = \frac{1}{\ln(r)} = \frac{1}{\ln(1.116123)} = \frac{1}{0.109861} = \$9.102392$$

$$u(x) = a - br^{-x} = a - b(1.116123)^{-x} = a - be^{-0.109861x} = a - be^{-x/9.102392}$$

Encoding an Exponential Utility Function using Risk Odds in ASSESS

This example uses the same data as the example above, i.e., -$$m$ = -\$10 and $$m$ = \$10, to illustrate how we can use the program ASSESS and successive bracketing to find the probability q=0.75 and the risk tolerance ρ = \$9.10 for a person with an exponential utility curve by encoding the utility for just one point (i.e., by setting \$0 as the Certain Equivalent for the lottery \$0 ~ (\$$m$, q; -\$$m$, 1-$q$)) using the "*Probability Equivalence*" Utility Assessment Method.



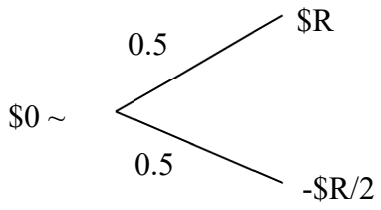
The final risk tolerance ρ = \$9.10 calculated by ASSESS for \$$m$ = \$10 and q = 0. 75 is the same as that calculated manually for the same values, \$$m$ and q, but with much less effort.

Assessing the Exponential Utility Function using Risk Tolerance

Using the parameter *risk tolerance* $\rho = 1/\gamma$ the exponential utility function can be written as:

$$u(x) = a - be^{-\frac{x}{\rho}} = a - b\exp(-x/\rho)$$

As shown below, the value R for which the following indifference statement holds true is approximately equal to the risk tolerance, ρ :



$$u(0) = \frac{1}{2}u(R) + \frac{1}{2}u(-\frac{R}{2})$$

$$a - be^{-\frac{0}{\rho}} = \frac{1}{2}(a - be^{-\frac{R}{\rho}}) + \frac{1}{2}(a - be^{+\frac{R}{2\rho}})$$

$$2 = e^{-\frac{R}{\rho}} + e^{+\frac{R}{2\rho}} = (e^{-\frac{R}{2\rho}})^2 + (e^{-\frac{R}{2\rho}})^{-1}$$

Let: $z = e^{-\frac{R}{2\rho}}$

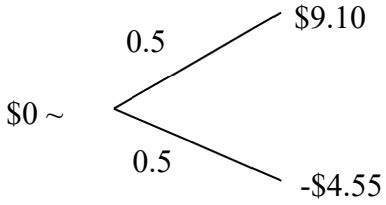
$$z^3 - 2z + 1 = 0$$

$$z = \{1;\ \frac{1}{2}(-1 \pm \sqrt{5})\} = \{-1.618;\ 1;\ 0.618\}$$

$$e^{-\frac{R}{2\rho}} = 0.618$$

$$\frac{R}{\rho} = -2\ln(0.618) = 0.962 \approx 1$$

For example, the risk tolerance $\rho = \$9.10$ found using the risk odds approach implies:



Risk Tolerance—Empirical Evidence

Ronald Howard[2], one of the founders of modern decision analysis, has discovered (through his consulting experience with large companies) tentative relationships between risk tolerance, ρ, and several financial variables, such as: net sales, net income, and equity.

Specifically, for the companies that Howard studied (in the oil and chemicals industries), risk tolerance ρ was found to be approximately:

- 6.4% of net sales,
- 124% of net income, and
- 15.7% of equity

Another study[3] has found that more generally and across industries risk tolerance ρ is about:

- 1/6 of equity, or
- 1/5 of market value

For example, according to these estimates, a company with equity equal to $30 million should have a risk tolerance of approximately ρ = $5 million.

Howard admits that these percentages are only guidelines. However, they do indicate that:

- Larger and more profitable companies tend to have larger values of ρ
- Which means that they are more willing to take risks involving large dollar amounts.

Advantages of Exponential Utility—Summary

The exponential utility curve has the computational advantages of the *delta property*: if all outcomes of a deal are increased by the same amount, Δ, then the certain equivalent of the deal will also increase by the amount Δ.

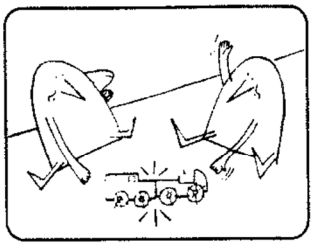
As a result:

1. There is no need to consider the initial wealth of the decision maker in evaluating uncertain deals. This is because the initial wealth can be adjusted up and down by any amount Δ.
2. When a delta-person is engaged in a sequence of independent lotteries, his certain equivalent for the whole prospect is the sum of the certain equivalents for each individual lottery.
3. The {EVPI | EVSI} about an uncertainty is the certain equivalent CE with free {perfect | sampled} information on this uncertainty less the CE without the information.
4. The personal indifferent [buying | selling] price, [PIBP | PISP], of any uncertain deal are equal to each other no matter whether a person owns the deal or not.
5. Risk aversion (and the shape of the exponential utility curve) is characterized by a single parameter: the risk aversion coefficient or its reciprocal, the risk tolerance.

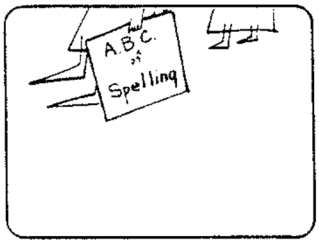
Of course, there are situations where the delta property does not hold and other utility curves, such as the logarithmic, need to be adopted.

[2] Howard, R. A. 1988. Decision Analysis: Practice and Promise. Management Science, 34, 679-695.
[3] McNamee, P., and Celona, J. 1990. Decision Analysis with Supertree, 2nd Ed, Scientific Press, South San Francisco, CA.

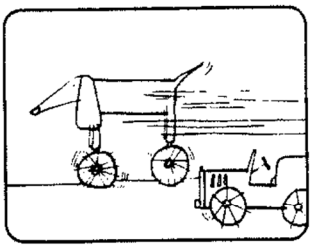
Jumping to Conclusions



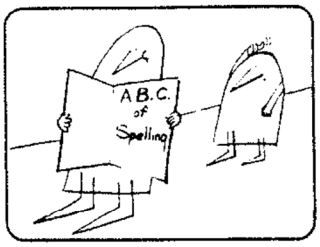
Statistics show that most car accidents occur when cars travel at moderate speeds, and that very few accidents occur at speeds of more than 150 kilometers per hour. Does this mean it is safer to drive at higher speeds?



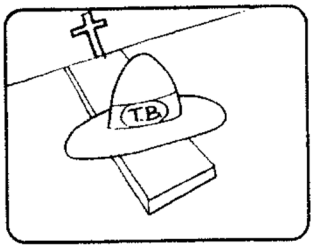
A research study showed that children with big feet could spell better than those with small feet. Does this mean that the size of one's foot is a measure of one's ability to spell?



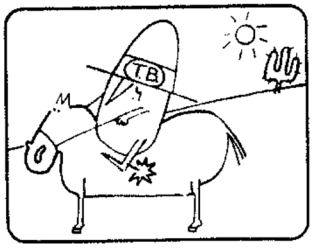
Not at all. Statistical relationships often have nothing to do with cause and effect. Most people drive at moderate speeds, so naturally most accidents occur at moderate speeds.



It does not. The study included *growing* children. All it showed was that older children, who, of course, have bigger feet, spell better than younger ones.



If statistics showed that more people died of tuberculosis in Arizona than in any other state, would that mean that Arizona's climate favored getting TB?



Quite the contrary. Arizona's climate is so *helpful* to TB victims that thousands of them go there. Naturally this would raise the average number of TB deaths.

© 2024 *Photios G. Ioannou*

These three episodes underscore the importance of not jumping to conclusions about cause and effect when you hear of a statistical correlation. Here are a few more instances:

1. It is often said that most car accidents occur near the home. Does this mean that travel on highways, many miles from home, is safer than driving around town? No. The statistics simply reflect the fact that a car is more often driven near one's home than on distant highways.

2. A study showed that a certain state ranked high both in the proportion of people who drink milk and in the proportion of people who die of cancer. Does this indicate that drinking milk causes cancer? No. The state also ranked high in the proportion of people who are elderly. Since cancer is commonly an affliction of the aged, this is what raised the proportion of cancer deaths.

3. A study showed that in a certain city there was a sharp rise both in deaths from heart failure and in the consumption of beer. Could it be that beer drinking increases the probability of a heart attack? No. The increase in both cases was the result of a rapid rise in population. By the same reasoning, heart attacks could be attributed to hundreds of other things: increased consumption of coffee, increased chewing of gum, increased playing of bridge, increased television watching.

4. A study showed that at the same time a certain European city had a large increase in population, there was a large increase in the number of stork nests in the city. Does this support the belief that storks bring babies? No, it reflected the fact that as the number of buildings increased, there were more places in the city where storks could nest.

5. A recent study showed that most great mathematicians were eldest sons. Does this mean that first-born sons are more likely to have mathematical ability than sons born later? No, it simply reflects the surprising fact that most sons are eldest sons.

The last example suggests some interesting experiments. Survey your male friends to see if more than half are eldest sons. Do the same with your female friends to see what proportion are eldest daughters.

Or a thought experiment. Consider a population in which 100 families each have two children. What fraction of the boys (or girls) will be eldest sons (or daughters)? (Answer: 3/4). Compute the fraction in a population of 100 families, each with three children. (Answer: 7/12.) It goes without saying that in families of only one child, that child will be the eldest.

The exact percentage of eldest children of one sex obviously will vary with the sizes of families in the population under consideration, but for all populations it is over 1/2 and for most populations the figure is well over 1/2.

These examples may stimulate you to look for other instances of statistical statements that are easily misinterpreted with respect to cause and effect. Modern advertising, especially television commercials, is a rich source of such misleading assertions.

© 2024 *Photios G. Ioannou*

Combinations of Risks — Total Project Cost

The total cost of a project is a function of its various cost items:

$$\text{Project Cost} = \sum_i \text{Cost}_i$$

$$E[\text{Project Cost}] = E[\sum_i \text{Cost}_i] = \sum_i E[\text{Cost}_i]$$

The expected value of the total project cost is the sum of the expected values of all *n* cost items $E[\text{Cost}_i]$ (this is always true). If the *n* cost items cost items are mutually independent, then the variance of the total project cost is given by:

$$Var[\text{Project Cost}] = Var[\sum_i \text{Cost}_i] = \sum_i Var[\text{Cost}_i]$$

In general, the *n* cost items are *not* independent and covariance terms must be considered:

$$Var[\text{Project Cost}] = Var[\sum_i \text{Cost}_i] = \sum_i \sum_j Cov[\text{Cost}_i, \text{Cost}_j]$$

$$= \sum_i Var[\text{Cost}_i] + 2\sum_{i=1}^{n}\sum_{j>i}^{n} Cov[\text{Cost}_i, \text{Cost}_j]$$

$$= \sum_i \sigma_i^2 + 2\sum_{i=1}^{n}\sum_{j>i}^{n} \rho_{ij}\sigma_i\sigma_j$$

The covariance between any two cost items (*i*,*j*) is $Cov[\text{Cost}_i, \text{Cost}_j] = \rho_{ij}\sigma_i\sigma_j$ where $\sigma_i^2 = Var[\text{Cost}_i]$, and ρ_{ij} is the correlation coefficient. The coefficient ρ_{ij} can be estimated subjectively according to the degree of correlation expected as follows:

Relative correlation	ρ
exact	1
strong	0.85
average	0.50
very small, if any	0.15
none; independent	0
exact inverse	-1

As an example, here is the simplest case that involves the variance of the *sum* or the *difference* of only *two* random variables *X* and *Y*:

$$Var[X \pm Y] = Var[X] + Var[Y] \pm 2Cov[X, Y]$$

$$\sigma_{X\pm Y}^2 = \sigma_X^2 + \sigma_Y^2 \pm 2\rho_{XY}\sigma_X\sigma_Y$$

Notice that the sign of the variance terms is always positive. The sign of the covariance terms $Cov[X, Y]$ or $\rho_{XY}\sigma_X\sigma_Y$ is given by (sign of *X*)*(sign of *Y*).

Range Estimate

A range estimate communicates not only an expected value but a measure of its variability, and is often expressed in the form of two numbers separated by a plus-and-minus sign: $m \pm \sigma$.

Example. A project has two items which are estimated to be particularly risky. These are (1) the cost of the foundation with possible dewatering problems and (2) the cost of demolishing the building which is now on the lot and recycling some of its materials as required by the contract.

The following costs are estimated and considered independent:

Cost of project except foundation and demolition	$m_1 = 1{,}600^k$; $\sigma_1 = 30^k$
Cost of foundation	$m_2 = 100^k$; $\sigma_2 = 15^k$
Cost of demolition	$m_3 = 40^k$; $\sigma_3 = 10^k$

The cost of the project can then be estimated as

$$m_{\text{Project Cost}} = m_1 + m_2 + m_3 = 1{,}600^k + 100^k + 40^k = 1{,}740^k$$

$$\sigma^2_{\text{Project Cost}} = \sigma_1^2 + \sigma_2^2 + \sigma_3^2 = 30^2 + 15^2 + 10^2 = 900 + 225 + 100 = 1225^{k^2}$$

$$\sigma_{\text{Project Cost}} = \sqrt{1225}^{\,k} = 35^k$$

Therefore, the range for the estimated project cost is $1740^k \pm 35^k$

Probability Distributions in Decision Trees

Decision trees work best, particularly with manual calculations (as opposed to computer calculations), using discrete outcomes (as opposed to continuous). However, many random variables are estimated or defined by continuous distributions. If a distribution can be considered bell shaped and not highly skewed (i.e., it is close to symmetrical), we can make the following assumption and substitute the PDF of a continuous chance variable with an "equivalent" discrete PMF that has almost the same mean and variance:

Outcome	P[Outcome]
$m - 2\sigma$	0.0833
$m - \sigma$	0.1667
m	0.5000
$m + \sigma$	0.1667
$m + 2\sigma$	0.0833

Notice that this approximation of a continuous PDF by a discrete PMF does not treat what are known as the tails of the continuous distribution (outcomes with very small probabilities) which may be important with some utility curves.

Example. In analyzing the example problem used earlier where project cost was found to be $1525^k \pm 39^k$, the following outcomes could be considered:

Cost	P[Cost]
$1525^k - 2(39^k) = 1447^k$	0.0833
$1525^k - (39^k) = 1486^k$	0.1667
1525^k	0.5000
$1525^k + (39^k) = 1564^k$	0.1667
$1525^k + 2(39^k) = 1603^k$	0.0833

We can then use these dollar outcomes and their respective probabilities to analyze a decision problem using a utility curve.

MOMENTS OF THE NORMAL DISTN

$X \sim N(m, \sigma^2)$ $\quad f_X(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-m}{\sigma}\right)^2}$

Moment Generating Function: $M(t) = e^{mt + \frac{1}{2}\sigma^2 t^2}$

$M'(t) = (m + \sigma^2 t) M(t)$

$M''(t) = (\sigma^2 + (m+\sigma^2 t)^2) M(t)$

$M'''(t) = (3\sigma^2(m+\sigma^2 t) + (m+\sigma^2 t)^3) M(t)$

$M^{(4)}(t) = (3\sigma^4 + 6(m+\sigma^2 t)^2\sigma^2 + (m+\sigma^2 t)^4) M(t)$

$M^{(5)}(t) = (15\sigma^4(m+\sigma^2 t) + 10\sigma^2(m+\sigma^2 t)^3 + (m+\sigma^2 t)^5) M(t)$

$M^{(6)}(t) = [15\sigma^6 + 45\sigma^4(m+\sigma^2 t)^2 + 15\sigma^2(m+\sigma^2 t)^4 + (m+\sigma^2 t)^6] M(t)$

$m = 0, \sigma = 1$

$E[X] = M'(0) = m = 0$

$E[X^2] = M''(0) = \sigma^2 + m^2 = \sigma^2 = 1$

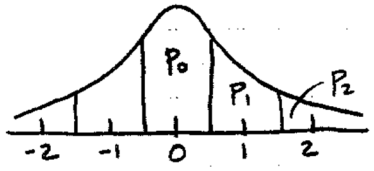
$E[X^3] = M'''(0) = m^3 + 3m\sigma^2 = 0$

$E[X^4] = M^{(4)}(0) = m^4 + 6m^2\sigma^2 + 3\sigma^4 = 3$

$E[X^5] = M^{(5)}(0) = m^5 + 10m^3\sigma^2 + 15m\sigma^4 = 0$

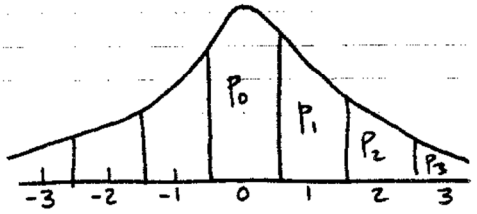
$E[X^6] = M^{(6)}(0) = m^6 + 15m^4\sigma^2 + 45m^2\sigma^4 + 15\sigma^6 = 15$

DISCRETIZING THE NORMAL DISTR. KEEPING MOMENTS SAME:



$$E[X^2] = (p_1 + 4p_2)2 = 1$$
$$E[X^4] = (p_1 + 16p_2)2 = 3 \Rightarrow$$

5 points: $p_0 = \frac{1}{2}$ $p_1 = \frac{1}{6}$ $p_2 = \frac{1}{12}$



$$E[X^2] = (p_1 + 4p_2 + 9p_3)2 = 1$$
$$E[X^4] = (p_1 + 16p_2 + 81p_3)2 = 3$$
$$E[X^6] = (p_1 + 64p_2 + 729p_3)2 = 15 \Rightarrow$$

7 points: $p_0 = \frac{7}{18}$ $p_1 = \frac{1}{4}$ $p_2 = \frac{1}{20}$ $p_3 = \frac{1}{180}$

Subjective Probability Assessment using PERT Formulas

The mean and variance of a random variable can be estimated subjectively using the formulas suggested by PERT based on three point estimates. For a cost item, we simply estimate the range of possible costs by estimating the lowest and highest possible costs, *a* and *b*, respectively, and the most likely outcome (the mode), *M*. The mean and standard deviation for this cost item are then given by the two PERT formulas:

$$m = \frac{a + 4M + b}{6}$$

$$\sigma = \frac{b - a}{6}$$

Example. One's best early estimate of cost for a structural frame is \$115^k. It is also estimated that the absolute minimum it could cost is \$95^k and the absolute maximum is \$160^k. The mean cost can be estimated by

$$m = \frac{95 + 4 \times 115 + 160}{6} = \frac{715}{6} = 119^k$$

The standard deviation would be estimated to be

$$\sigma = \frac{160 - 95}{6} = \frac{65}{6} = \$10.8^k$$

The probability distribution which is defined by these relationships is the beta distribution. The three-point estimates *a*, *M*, and *b* are shown on a beta distribution PDF in Fig. 1. The value of *M* is not midway between *a* and *b* and the distribution is skewed. Therefore, the most likely estimate, *M*, is not equal to the mean, *m*. The great variety of shapes that the beta distribution can take is shown in Fig. 2. The distributions are standardized with a range of one between *a* = 0 and *b* = 1. The shape of each PDF and CDF is shown to be directly related to the relative value of the most likely estimate within the range.

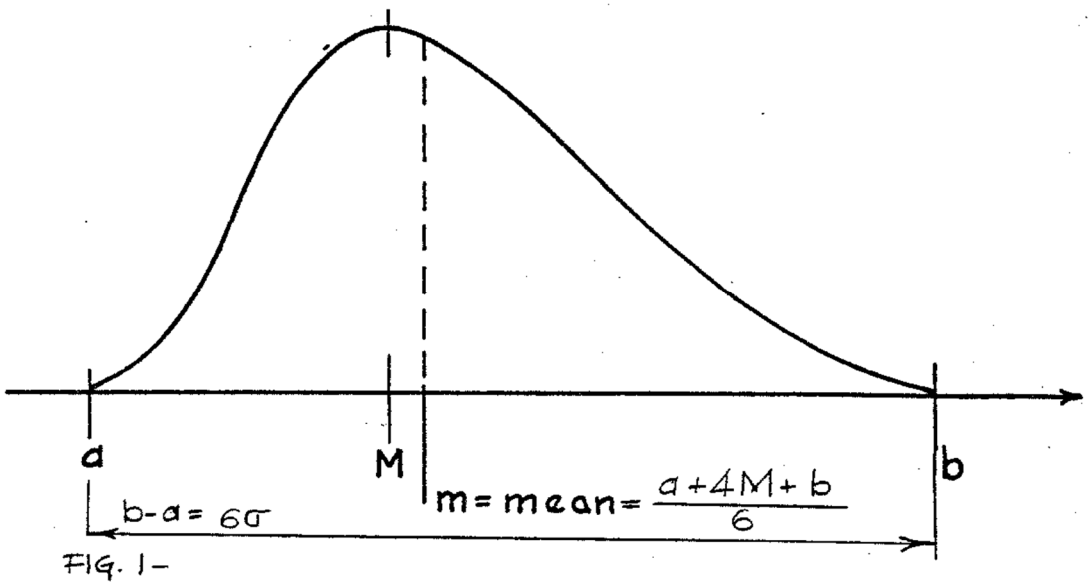


FIG. 1–

The most common estimates of time, cost, or whatever, are single values or point estimates rather than distributions. These individual estimates tend to be most likely estimates, modes, rather than expected values, means. These modes are then summed to get a total estimate of cost, time, or other unit, again a single value or point estimate.

We know that the expected value of a sum of random variables is the sum of their expected values. That is,

$$E[\sum_{i=1}^{n} X_i] = \sum_{i=1}^{n} E[X_i]$$

In general, however, the most likely value of a sum does not equal the sum of the most likely values of its components.

$$M[\sum_{i=1}^{n} X_i] \neq \sum_{i=1}^{n} M[X_i]$$

In fact, the sum of the modes of several random variables does not provide any meaningful quantity at all. One exception, of course, is when all random variables have symmetric distributions, in which case the modes are also the means. The other exception is when all random variables are perfectly correlated (i.e., they are linear functions of each other).

In conclusion, care should be taken that we sum the means and not the modes.

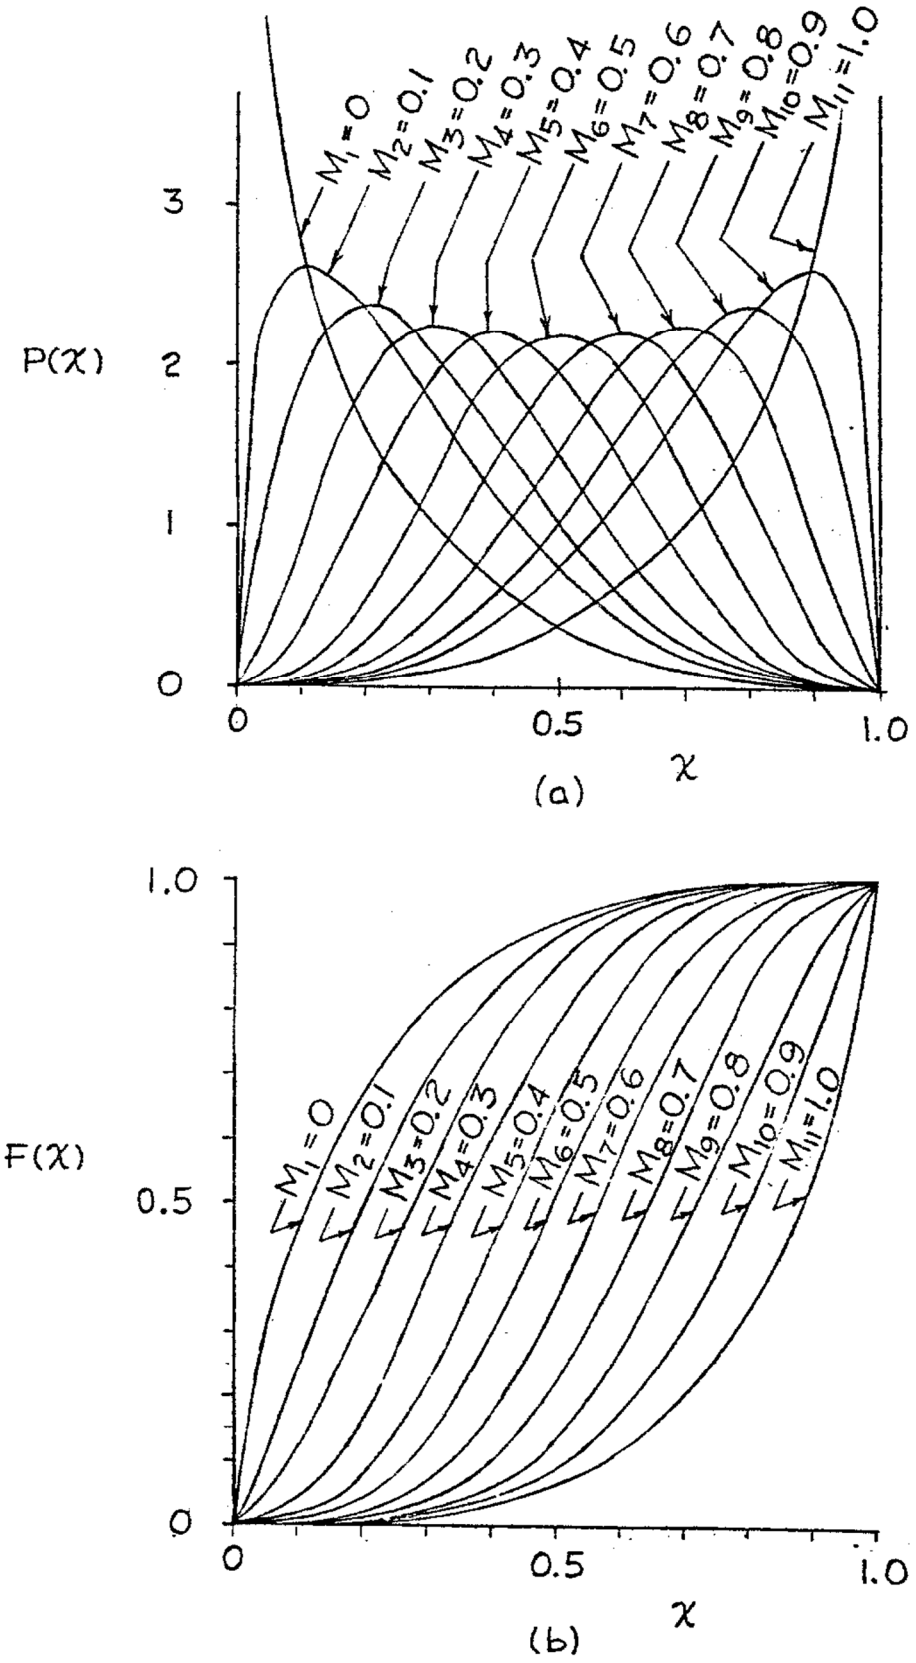


FIG. 2 - BETA DISTRIBUTIONS WITH DIFFERENT VALUES OF M

Subjective Probability Assessment using Perry and Greig Formulas

There are three problems in the use of PERT assumptions.

1. It is difficult to estimate the limits of a random variable, because so often at the tails anything seems possible.
2. There is no reason to consider all continuous distributions to be beta.
3. The formulas are not as accurate as might be desired.

Of these problems, only the first is of much concern to our estimation of outcomes and probabilities. Nevertheless, it is quite significant. The estimated variance is very important in making decisions and its accuracy is very sensitive to the estimated limits of a distribution. However, the probability of observing outcomes in the tails of a distribution close to the upper and lower limit is quite small. Thus, estimating the limits, *a* and *b*, might be very difficult.

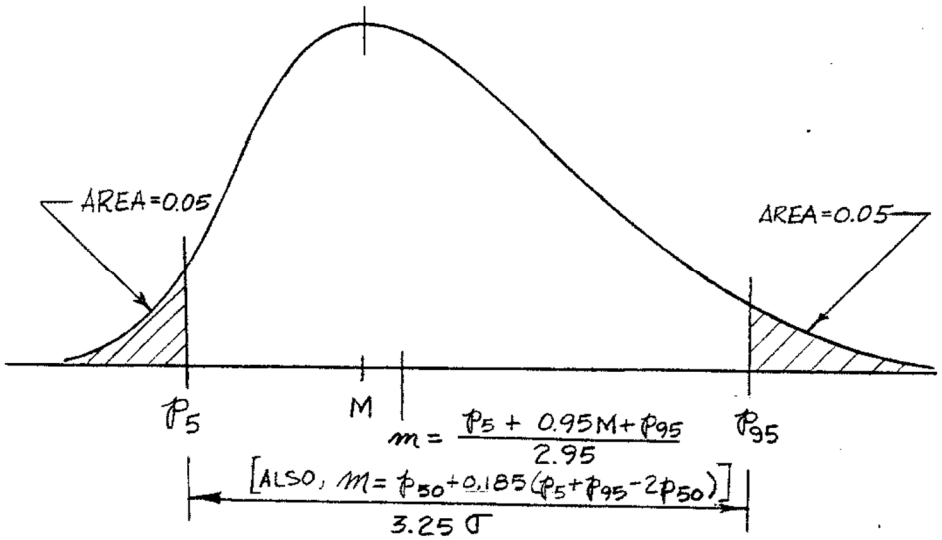
The suggested alternative is to use estimates of the 5th and 95th percentiles (p_5 and p_{95}), for which there is a 5% probability of being outside. The Perry and Greig formulas[4] use these percentiles, along with estimates of either the mode (M) or the 50th percentile (p_{50}) (i.e., the median) to estimate m and σ:

$$m = \frac{p_5 + 0.95M + p_{95}}{2.95}$$

$$m = p_{50} + 0.185(p_5 + p_{95} - 2p_{50})$$

$$= 0.63p_{50} + 0.185(p_5 + p_{95})$$

$$\sigma = \frac{p_{95} - p_5}{3.25}$$



Note that the percentiles, p_5, p_{50} and p_{95}, can be estimated easily using simple lotteries and a *probability wheel*. (In contrast, the extreme values, *a* and *b* (i.e., p_0 and p_{100}), used by PERT

[4] Perry, C., and Greig, I.D., "*Estimating the Mean and Variance of Subjective Distributions in PERT and Decision Analysis*," Management Science, Vol. 21, No. 12, Aug., 1975, pp. 1477-80.

cannot be estimated easily.) Moreover, Perry and Greig have shown that the above formulas are distribution-free — they are accurate for beta, gamma, lognormal, and normal distributions.

Because the utilities of the percentiles for money, give the corresponding percentiles for utility, i.e., $u(p_5) = u_5$, $u(p_{50}) = u_{50}$ and $u(p_{95}) = u_{95}$, Perry and Greig also suggest that we can use the last expression above to calculate directly the expected utility for a continuous lottery:

$$\begin{aligned} E[u] &= u(p_{50}) + 0.185[u(p_5) + u(p_{95}) - 2u(p_{50})] \\ &= 0.63u(p_{50}) + 0.185[u(p_5) + u(p_{95})] \end{aligned}$$

This is an extremely useful and accurate result. Although it considers a spread of only 3.25σ (versus the 4σ considered by the discretization suggested earlier), it gives greater weight (0.185) to the extremes, p_5 and p_{95}. It is also more applicable for skewed (asymmetrical) distributions.

Example. Determine the expected utility for Contractor I for a random variable of profit estimated as p_5 = -$50,000, M = $100,000 and p_{95} = $400,000. We can calculate the mean and standard deviation as follows,

$$m = \frac{-50,000 + 0.95(100,000) + 400,000}{2.95} = \$150,900$$

$$\sigma = \frac{400,000 - (-50,000)}{3.25} = \$138,500$$

Profit could be described as $150,900 ± $138,500.

We can calculate the expected the utility associated with the random variable profit by first calculating the median p_{50} as follows,

$$\begin{aligned} p_{50} &= \frac{m - 0.185(p_5 + p_{95})}{0.63} \\ &= \frac{150,900 - 0.185(-50,000 + 400,00)}{0.63} = \$136,800 \end{aligned}$$

Thus,

$$\begin{aligned} E[u] &= 0.63u_{50} + 0.185[u_5 + u_{95}] \\ &= 0.63u(p_{50}) + 0.185[u(p_5) + u(p_{95})] \\ &= 0.63u(136,800) + 0.185[u(-50,000) + u(400,000)] \\ &= 0.63(81) + 0.185[63 + 92.3] \\ &= 79.8 \end{aligned}$$

Notice that the Perry and Greig formulas allowed us to calculate the expected utility of a continuous random variable (profit) without the need to approximate the continuous PDF of this variable by a discrete PMF as is necessary when using PERT formulas. Moreover, we did not have to make any assumptions about the shape or type of the profit PDF. This is contrast to the PERT formulas that assume that the random variable follows a beta distribution.

Expected Utility of a Continuous Random Variable

A simple approach for finding the expected utility of a continuous random variable is to approximate the continuous monetary and utility distributions by making them discrete using the 5-branch approximation presented earlier to get $u(m)$, $u(m \pm \sigma)$ and $u(m \pm 2\sigma)$ with probabilities of 1/2, 1/6, and 1/12, respectively.

x_i (\$)	$u_i = u(x_i)$	$P[X = x_i] = P[U = u_i]$	$P[X = x_i] = P[U = u_i]$
$m - 2\sigma$	$u(m - 2\sigma)$	0.0833	1/12
$m - \sigma$	$u(m - \sigma)$	0.1667	1/6
m	$u(m)$	0.5000	1/2
$m + \sigma$	$u(m + \sigma)$	0.1667	1/6
$m + 2\sigma$	$u(m + 2\sigma)$	0.0833	1/12

For a discrete monetary lottery (distribution), such as the approximation above, the probabilities of the monetary outcomes, x_i, are also the probabilities of the utility of those outcomes, $u_i = u(x_i)$. Thus, it is simple to compute the expected utility of the lottery:

$$EUV = E[u(X)] = \sum_i u(x_i) P[X = x_i]$$

Perry and Greig Formulas for the Expected Utility of a Continuous Random Variable

A more accurate approach is to approximate the continuous distribution for utility using percentiles. Given the percentiles for the monetary distribution, p_5, p_{50}, and p_{95}, the percentiles for the utility distribution are $u_5 = u(p_5)$, $u_{50} = u(p_{50})$, $u_{95} = u(p_{95})$:

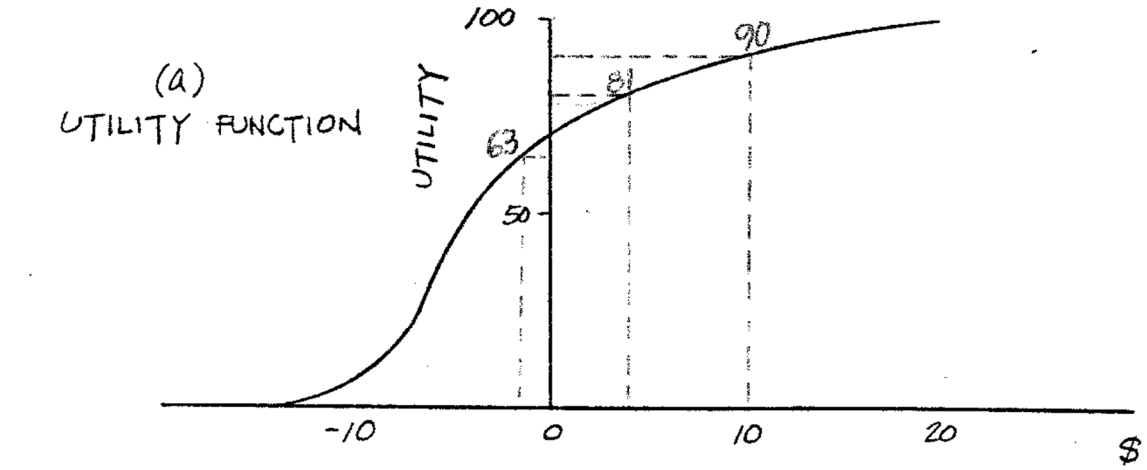
x_i (\$)	$u_i = u(x_i)$	$P[X = x_i] = P[U = u_i]$
p_5	$u_5 = u(p_5)$	0.185
p_{50}	$u_{50} = u(p_{50})$	0.630
p_{95}	$U_{95} = u(p_{95})$	0.185

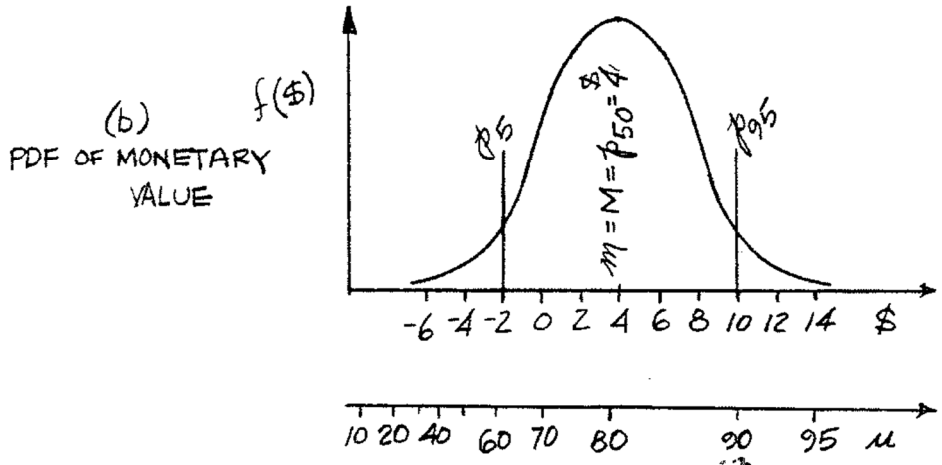
From these we can calculate expected utility by the Perry and Greig formula:

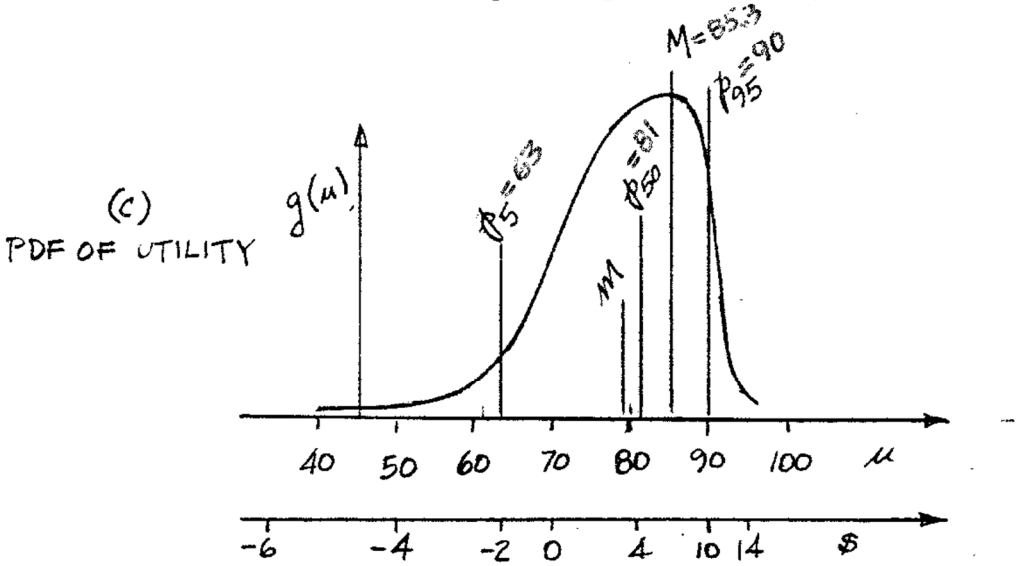
$$E[u(X)] = 0.185u_5 + 0.63u_{50} + 0.185u_{95}$$

Example. The figure on the following page shows the relationship between the utility curve and the monetary and utility PDFs. The utility curve is shown in (a). The monetary probability distribution is shown in (b). It is symmetrical and thus mode (M) = mean (m) = median (p_{50}). Below the monetary scale of the monetary PDF in (b) is shown a utility scale for the same PDF. The utility values in this scale correspond to the monetary values on the monetary scale and as a result they are not equally spaced. If they are made equally spaced, we get the utility PDF of (c). Because the utility function is not linear, the shape of the utility PDF is different from that of the monetary PDF. Also, its mean and mode have changed. However, the correspondence between monetary and utility percentiles has not changed: $u_5 = u(p_5)$, $u_{50} = u(p_{50})$, $u_{95} = u(p_{95})$.

At the bottom of the figure are calculations for the mean and mode of the utility distribution based on the percentiles of utility, $u_5 = u(p_5)$, $u_{50} = u(p_{50})$, $u_{95} = u(p_{95})$, which are simply the utilities of the monetary distribution percentiles.







$$E(u) = m_u = .630\,p_{50} + .185(p_5 + p_{95}) = .630(81) + .185(63 + 90) = 79.3$$

$$M_u = \frac{2.95m - p_5 - p_{95}}{0.95} = \frac{2.95(79.3) - 63 - 90}{0.95} = 85.3$$

The Distribution of a Utility Function

Let X be a random variable of interest (e.g., $ profit)

Let U be the utility function for X, $U = g(x)$ (e.g., the utility for $ profit)

Since X is a random variable and U is a function of X, then U is also a random variable. Because U is a monotonically increasing function of X, we can find the CDF of U as follows:

$$F_U(u) = P[U \le u] = P[X \text{ takes on such values such that } g(x) \le u] = P[X \le g^{-1}(u)] = F_X(g^{-1}(u))$$

The PDF of U is then found be differentiation as follows:

$$f_U(u) = \frac{F_U(u)}{du} = \frac{F_X(x)}{dx}\frac{dx}{du} = f_X(x)\frac{dg^{-1}(u)}{du} = f_X(g^{-1}(u))\left|\frac{dg^{-1}(u)}{du}\right|$$

Notice that in the above formula, the absolute value addresses situations where the utility function is a monotonically decreasing (rather than increasing) function of X, such as when larger values of X represent larger costs.

Example

Let $U = g(x) = \sqrt{x} \quad x \ge 0$ and X follows an exponential distribution, $f_X(x) = \lambda e^{-\lambda x} \quad x \ge 0$.

From the above we have that, $X = g^{-1}(u) = u^2 \quad u \ge 0$, and $F_X(x) = 1 - e^{-\lambda x} \quad x \ge 0$

Thus,

$$F_U(u) = F_X(g^{-1}(u)) = 1 - e^{-\lambda u^2} \quad u \ge 0$$

$$f_U(u) = \frac{F_U(u)}{du} = 2\lambda u e^{-\lambda u^2} \quad u \ge 0$$

We can also find the same result directly by applying the above formula for the PDF of U.

$$f_U(u) = f_X(g^{-1}(u))\left|\frac{dg^{-1}(u)}{du}\right| = 2\lambda u e^{-\lambda u^2} \quad u \ge 0$$

There are two approaches to finding the expected utility. One is to use $f_U(u)$ and find the expected value of U:

$$E[U] = \int_0^\infty u f_U(u)du = \int_0^\infty u\, 2\lambda u e^{-\lambda u^2} du = \int_0^\infty 2\lambda u^2 e^{-\lambda u^2} du = \frac{1}{2}\sqrt{\frac{\pi}{\lambda}}$$

The other approach is to use $f_X(x)$ and find the expected value of $g(x)$:

$$E[U] = E[g(x)] = \int_0^\infty g(x) f_X(x)dx = \int_0^\infty x^{1/2}\lambda e^{-\lambda x} dx = \frac{1}{2}\sqrt{\frac{\pi}{\lambda}}$$

The certain equivalent $CE[X]=\tilde{x}$ for this lottery satisfies the definition

$$u(\tilde{x})=E[U]=E[u(x)]$$

So,

$$\sqrt{\tilde{x}}=\frac{1}{2}\sqrt{\frac{\pi}{\lambda}}$$

$$CE[X]=\tilde{x}=\frac{\pi}{4\lambda}=\frac{0.7854}{\lambda}$$

The risk premium $RP[X]=x_P$ for this lottery is:

$$RP[X]=E[X]-CE[X]=m_X-\tilde{x}=\frac{1}{\lambda}-\frac{\pi}{4\lambda}=\frac{1}{\lambda}(1-\frac{\pi}{4})=\frac{0.2146}{\lambda}$$

Example

Let $U=g(x)=-\mathrm{e}^{-x/\rho}\quad x\geq 0$. X follows an exponential distribution, $f_X(x)=\lambda e^{-\lambda x}\quad x\geq 0$.

From the above we have that, $X=g^{-1}(u)=-\rho\ln(-u), -1\leq u\leq 0$, and $F_X(x)=1-e^{-\lambda x}, x\geq 0$.

Thus,

$$F_U(u)=F_X(g^{-1}(u))=1-e^{\lambda\rho\ln(-u)}=1-(-u)^{\lambda\rho}, -1\leq u\leq 0$$

$$f_U(u)=\frac{F_U(u)}{du}=\lambda\rho(-u)^{\lambda\rho-1}, -1\leq u\leq 0$$

We can also find the same result directly by applying the above formula for the PDF of *U*.

$$f_U(u)=f_X(g^{-1}(u))\left|\frac{dg^{-1}(u)}{du}\right|=\lambda\rho(-u)^{\lambda\rho-1}, -1\leq u\leq 0$$

There are two approaches to finding the expected utility. One is to use $f_U(u)$ and find the expected value of *U*:

$$E[U]=\int_{-1}^{0}uf_U(u)du=\int_{-1}^{0}u\lambda\rho(-u)^{\lambda\rho-1}du=\int_{-1}^{0}-\lambda\rho(-u)^{\lambda\rho}du=\frac{\lambda\rho}{\lambda\rho+1}\left[(-u)^{\lambda\rho+1}\right]_{-1}^{0}=-\frac{\lambda\rho}{\lambda\rho+1}$$

The other approach is to use $f_X(x)$ and find the expected value of $g(x)$:

$$E[U]=E[g(x)]=\int_{0}^{\infty}g(x)f_X(x)dx=\int_{0}^{\infty}-\mathrm{e}^{-x/\rho}\,\lambda e^{-\lambda x}dx=\int_{0}^{\infty}-\lambda\,\mathrm{e}^{-x(\lambda+1/\rho)}\,dx=-\frac{\lambda\rho}{\lambda\rho+1}$$

The certain equivalent $CE[X]=\tilde{x}$ for this lottery satisfies the definition

$$u(\tilde{x}) = E[U] = E[u(x)]$$

So,

$$-\mathrm{e}^{-\tilde{x}/\rho} = -\frac{\lambda\rho}{\lambda\rho+1}$$

$$CE[X] = \tilde{x} = -\rho \ln(\frac{\lambda\rho}{\lambda\rho+1})$$

The risk premium $RP[X] = x_P$ for this lottery is:

$$RP[X] = E[X] - CE[X] = m_X - \tilde{x} = \frac{1}{\lambda} + \rho \ln(\frac{\lambda\rho}{\lambda\rho+1})$$

Does the Mode of X Correspond to the Mode of U?

The question we would like to address is:

$$u(M_X) = M_U \, ?$$

$$f_U(u) = f_X(g^{-1}(u)) \frac{dg^{-1}(u)}{du}$$

$$\frac{df_U(u)}{du} = \frac{df_X(x)}{dx}\left(\frac{dg^{-1}(u)}{du}\right)^2 + f_X(x)\frac{d^2g^{-1}(u)}{du^2}$$

So, the original question becomes:

$$\left.\frac{df_U(u)}{du}\right|_{u(M_X)} = \left.\frac{df_X(x)}{dx}\right|_{M_X}\left(\frac{dg^{-1}(u)}{du}\right)^2 + f_X(M_X)\left.\frac{d^2g^{-1}(u)}{du^2}\right|_{u(M_X)}$$

$$= f_X(M_X)\left.\frac{d^2g^{-1}(u)}{du^2}\right|_{u(M_X)} = 0\,?$$

Since $f_X(M_X)$ is the peak of $f_X(x)$ it cannot be zero. Thus, for the above to be true we must investigate the conditions for the following to be true:

$$\left.\frac{d^2g^{-1}(u)}{du^2}\right|_{u(M_X)} = 0\,?$$

This is true for a linear utility function (risk-neutral behavior). However, for a risk-neutral person, there is no need for a utility function. It is best to use expected monetary values.

For consistently risk-averse behavior the above condition is not true because risk aversion requires that:

$$\frac{d^2u}{dx^2} < 0 \quad \Rightarrow \quad \frac{d^2x}{du^2} < 0$$

Conclusion: The *utility of the mode of profit* does not give the *mode of the utility of profit*.

$$u(M_X) \neq M_U$$

For this reason, we cannot use the Perry and Greig formula that uses *the mode of utility* to estimate the *expected utility of profit*. Instead, to estimate the *expected utility of profit*, we must use the Perry and Greig formula that uses *the three percentiles of utility*:

$$E[u] = 0.63u_{50} + 0.185[u_5 + u_{95}]$$

The percentiles of utility in the above Perry & Greig formula are given by the utility of the corresponding percentiles of profit:

$$u_5 = u(p_5), \quad u_{50} = u(p_{50}), \quad u_{95} = u(p_{95})$$

Conditional Distributions

The conditional CDF of X is:

$$F_{X|a\le X\le b}(x)=P[X\le x\,|\,a\le X\le b]=\frac{P[X\le x\cap a\le X\le b]}{P[a\le X\le b]}=\frac{F_X(x)-F_X(a)}{F_X(b)-F_X(a)},\quad \text{for } a\le x\le b$$

The conditional PDF of X is:

$$f_{X|a\le X\le b}(x)=\frac{dF_{X|a\le X\le b}(x)}{dx}=\frac{f_X(x)}{F_X(b)-F_X(a)},\quad \text{for } a\le x\le b$$

Conditional Expected Value

The conditional expected value of X is

$$E[X\,|\,a\le X\le b]=\int_a^b xf_{X|a\le X\le b}(x)dx=\frac{1}{F_X(b)-F_X(a)}\int_a^b xf_X(x)dx$$

Example: The conditional expected value of the standard normal variate $U\sim N(0,1)$:

$$f_U(u)=\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}u^2}$$

The static moment of the area under the density function from *a* to *b* is:

$$\int_a^b uf_U(u)du=\int_a^b u\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}u^2}du=-\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}u^2}\Big|_a^b=-f_U(u)\Big|_a^b=f_U(a)-f_U(b)$$

$$E[U\,|\,a\le U\le b]=\frac{f_U(a)-f_U(b)}{F_U(b)-F_U(a)}$$

Example: The conditional expected value of the normal variate $X\sim N(m_X,\sigma_X^2)$:

$$X=m_X+\sigma_X U$$

$$E[X\,|\,a\le X\le b]=E[m_X+\sigma_X U\,|\,a\le(m_X+\sigma_X U)\le b]$$

$$=m_X+\sigma_X E[U\,|\,\frac{a-m_X}{\sigma_X}\le U\le\frac{b-m_X}{\sigma_X}]$$

$$=m_X+\sigma_X\frac{f_U(\frac{a-m_X}{\sigma_X})-f_U(\frac{b-m_X}{\sigma_X})}{F_U(\frac{b-m_X}{\sigma_X})-F_U(\frac{a-m_X}{\sigma_X})}$$

$$E[X\,|\,a\le X\le b]=m_X+\sigma_X^2\frac{f_X(a)-f_X(b)}{F_X(b)-F_X(a)}$$

Discretizing Continuous Probability Distributions

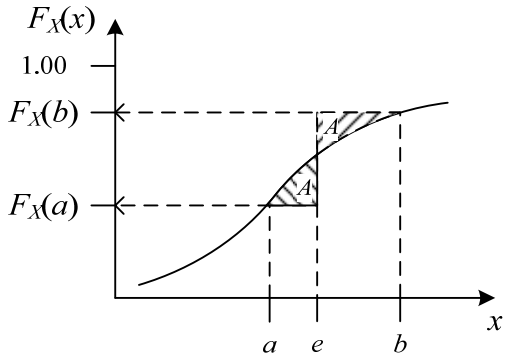
A typical procedure for discretizing continuous distributions is to divide the continuous CDF into a number of segments (typically three or five) and to choose a value in each segment to represent that segment with corresponding probability.

For example, all outcomes of the random variable X in the segment from a to b are to be represented by the single value e, where $a < e < b$. The probability of the single outcome e is then $F(b)$-$F(a)$. The problem is how to choose the appropriate value e to represent the interval a to b. A popular approach is to choose e to be the conditional expected value of X, given that $a < X < b$. An easy way to find e is to select it graphically based on the following theorem.

Theorem: If e is that value of x that makes the two shaded regions shown below have equal areas A, then e equals the conditional expected value of X, given that $a < X < b$:

$$e = E[X \mid a \le X \le b]$$

Special case: When a = -inf and b = +inf then e is the mean or expected value of X.



Proof: Let us assume that the left area A equals the right area A. Working from here we have:

$$\int_a^e F_X(x)dx - (e-a)F_X(a) = (b-e)F_X(b) - \int_e^b F_X(x)dx$$

$$\int_a^b F_X(x)dx = bF_X(b) - aF_X(a) - [F_X(b) - F_X(a)]e$$

$$e = \frac{1}{F_X(b) - F_X(a)}\{bF_X(b) - aF_X(a) - \int_a^b F_X(x)dx\}$$

Using integration by parts we can show that the term in curly braces is given by:

$$\int_a^b xf_X(x)dx = [xF_X(x)]_a^b - \int_a^b F_X(x)dx = bF_X(b) - aF_X(a) - \int_a^b F_X(x)dx$$

Thus, e is equals the definition of conditional expected value of X, given that $a < X < b$:

$$e = \frac{1}{F_X(b) - F_X(a)}\int_a^b xf_X(x)dx = E[X \mid a \le X \le b]$$

SUBJECTIVE PROBABILITY ASSESSMENT

MODELING VS. ENCODING

- D.A. DISTINGUISHES BETWEEN **DECISION** AND **STATE** VARIABLES
  - IF VARIABLES ARE NOT CLEAR, REFINE VARIABLE DEFINITIONS SO AS TO DISTINGUISH BETWEEN
    - **CONTROLLABLE** (DECISION), AND
    - **UNCONTROLLABLE** (STATE) VARIABLES
- BALANCE MODELING AND ENCODING
  - PROBLEM:
    - ENCODE PROBABILITIES AT PRESENT LEVEL OF DETAIL, OR
    - MODEL THE PROBLEM FURTHER AND ENCODE PROBABILITIES AT A FINER LEVEL OF DETAIL?
  - ANSWER:
    - CONSIDER THE IMPORTANCE OF THE DECISION PROBLEM
    - REFINE GROSS MODEL UNTIL
      - MODEL BENEFITS (IMPROVEMENT IN INFORMATION) < COSTS OF MODELING

GUIDELINES FOR PROBABILITY ENCODING

- ENCODE ONLY VARIABLES IMPORTANT TO THE DECISION AS DETERMINED BY A SENSITIVITY ANALYSIS
- DEFINE THE QUANTITY AS AN UNAMBIGUOUS RANDOM (UNCONTROLLABLE, STATE) VARIABLE
- STRUCTURE THE QUANTITY CAREFULLY. INCLUDE DEPENDENCE ON OTHER RANDOM VARIABLES AND STATE VARIABLES EXPLICITLY
- DEFINE THE QUANTITY CAREFULLY TO PASS THE CLAIRVOYANT TEST. CAN AN ANSWER BE GIVEN WITHOUT PROVIDING FURTHER CLARIFICATION?
- USE A SCALE OF MEASUREMENT THAT IS MEANINGFUL TO THE HUMAN SUBJECT (EXPERT): MILES VS KM, KG-FORCE VS NEWTON

ASSESSMENT CUES

- FEW PEOPLE REALIZE THE CUES ON WHICH JUDGMENT IS BASED.
  - E.G., HAZINESS IS CUE FOR ESTIMATING DISTANCE
  - THESE CUES INTRODUCE BIAS IN THE ENCODING PROCESS
- IT IS DIFFICULT TO CONTROL THE CUES THAT PEOPLE USE.
  - HAZY OBJECTS APPEAR MORE DISTANT EVEN IF WE KNOW WHY
- PEOPLE CAN BE MADE AWARE OF CUE BIASES AND CORRECT THEM.
  - PILOTS CAN BE TRAINED TO COMPENSATE FOR HAZINESS

© 2024 *Photios G. Ioannou*

BIASES IN PROBABILITY ENCODING

- BIAS: CONSCIOUS OR SUBCONSCIOUS DISCREPANCIES BETWEEN
  - THE SUBJECT'S RESPONSES, AND
  - AN ACCURATE DESCRIPTION OF THE SUBJECT'S UNDERLYING KNOWLEDGE
- DISPLACEMENT BIAS
  - SHIFTS THE WHOLE DISTRIBUTION UP OR DOWN
- VARIABILITY BIAS
  - CHANGE IN THE DISTRIBUTION'S SHAPE
  - USUALLY CENTRAL BIAS (I.E., OVERCONFIDENCE)

CAUSES OF BIAS

- MOTIVATIONAL BIASES
  - CONSCIOUS OR SUBCONSCIOUS ADJUSTMENTS IN THE SUBJECT'S RESPONSES MOTIVATED BY
    - THE REWARD SYSTEM
    - THE PERCEPTION OF THE REWARD SYSTEM
  - EXAMPLES:
    - WHEN EXCEEDED, A LOW SALES FORECAST SHOULD MAKE THE SALESMAN LOOK GOOD
    - ESTIMATOR MAY UNDERESTIMATE COSTS TO INCREASE CHANCE OF WINNING PROJECT AND THUS PLEASE BOSS
    - PROJECT MANAGER OVERESTIMATES COSTS (TIME) TO MAKE SURE PROJECT IS WITHIN BUDGET (SCHEDULE)
- COGNITIVE BIASES
  - CONSCIOUS OR SUBCONSCIOUS ADJUSTMENTS IN THE SUBJECT'S RESPONSES THAT ARE
    - SYSTEMATICALLY INTRODUCED BY THE WAY THE SUBJECT PROCESSES HIS/HER PERCEPTIONS
  - DEPEND ON MODES OF JUDGMENT

BASIC MODES OF JUDGMENT

- AVAILABILITY
  - THE EASE WITH WHICH INFORMATION IS RECALLED OR VISUALIZED
  - EXAMPLES:
    - EVENTS THAT CAUSED STRONG IMPRESSIONS
    - RECENT EVENTS
  - AVAILABILITY CAN BE USED TO COMPENSATE FOR CENTRAL BIAS BY DELIBERATELY ASKING FOR SCENARIA OF EXTREME EVENTS
- ADJUSTMENT AND ANCHORING
  - THE MOST READILY AVAILABLE PIECE OF INFORMATION OFTEN FORMS AN INITIAL BASIS FOR FORMULATING RESPONSES
  - SUBSEQUENT RESPONSES ARE ADJUSTMENTS FROM THIS BASIS
  - EXTREME OUTCOMES SHOULD BE ENCODED FIRST TO AVOID ANCHORING TO THE MEDIAN
- REPRESENTATIVENESS (JUDGMENT OF SIMILARITY)
  - THE PROBABILITY OF AN EVENT IS EVALUATED ACCORDING TO THE DEGREE THAT THE EVENT (SAMPLE) IS CONSIDERED TO BE REPRESENTATIVE OF, OR SIMILAR TO, SOME MAJOR CHARACTERISTIC OF THE UNDERLYING PROCESS OR POPULATION
  - ASSESSED DISTRIBUTION OF AVERAGE IS OFTEN SIMILAR TO DISTRIBUTION OF SAMPLE (OPPOSITE OF CENTRAL BIAS)
    - EXPERT USES UNDERLYING SAMPLE DISTRIBUTION AND IGNORES SAMPLE SIZE AND THE DISTR. OF THE AVERAGE
  - SINGLE PIECE OF “TEST” INFORMATION MAY BE CONSIDERED MORE REPRESENTATIVE THAN PAST “GENERAL” INFO.
    - EXPERT RELIES HEAVILY ON SINGLE TEST OUTCOME AND IGNORES EXTENSIVE PRIOR INFORMATION AS TOO “GENERAL” (I.E., VIOLATES BAYES THEOREM)
    - EXPLICITLY USE BAYES THEOREM WHENEVER POSSIBLE
- UNSTATED ASSUMPTIONS
  - ENCODED PROBABILITIES MAY BE CONDITIONAL ON UNSTATED ASSUMPTIONS
    - ASSUMPTIONS MUST BE EXPLICITLY STATED
    - ADDITIONAL MARGINAL DISTRIBUTIONS MUST BE ENCODED IF ASSUMPTIONS ARE NOT CERTAIN
- COHERENCE
  - ASSESSED PROBABILITIES MAY DEPEND ON THE EASE OF BUILDING SCENARIOS THAT LEAD TO SPECIFIC EVENT OR OUTCOME
  - PROBLEM IS SIMILAR TO CIRCUMSTANTIAL EVIDENCE. IT MAKES PEOPLE ASSIGN HIGHER PROBABILITIES TO EVENTS FOR WHICH A SCENARIO CAN BE EASILY CONSTRUCTED.

- ONE SHOULD DELIBERATELY ASK FOR SCENARIA THAT LEAD TO EXTREME EVENTS AND BALANCE THESE OUT.

ENCODING METHODOLOGIES

- P METHODS
  - SUBJECT ASSIGNS CDF PROBABILITIES TO GIVEN OUTCOMES
- V METHODS
  - SUBJECT ASSIGNS OUTCOMES TO GIVEN CDF PROBABILITIES
- PV METHODS
  - SUBJECT ASSIGNS POINTS ON THE CDF DIRECTLY

RESPONSE METHODS

- DIRECT RESPONSE METHODS
  - SUBJECT GIVES NUMBER DIRECTLY (PROB. OR OUTCOME)
- INDIRECT RESPONSE METHODS
  - SUBJECT CHOOSES BETWEEN LOTTERIES UNTIL A POINT OF INDIFFERENCE IS REACHED.

INDIRECT RESPONSE TECHNIQUES

- PROBABILITY WHEEL
  - BEST ENCODING TOOL
  - MOST EFFECTIVE WHEN USED AS A P-METHOD
  - MOST USEFUL FOR THE RANGE 0.1 TO 0.9
- FIXED PROBABILITY EVENTS (V METHODS)
  - PROB[10 HEADS IN A ROW] = 1/1000
  - PROB[ROYAL FLUSH] = 1/65,000

INTERVAL TECHNIQUE (V METHOD)

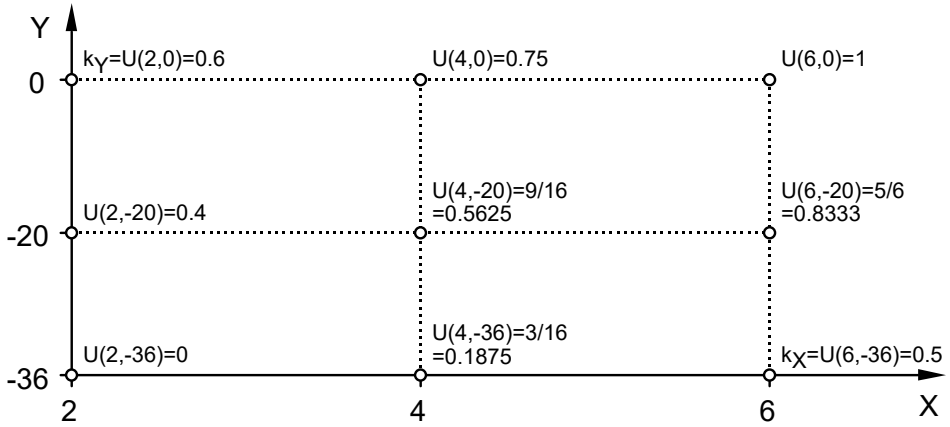
- TO FIND THE MEDIAN, DIVIDE THE RANGE OF OUTCOMES IN TWO UNTIL THEY ARE EQUALLY LIKELY
- TO FIND THE QUARTILES SUBDIVIDE EACH OF THE ABOVE INTERVALS IN TWO UNTIL THEY ARE EQUALLY LIKELY
- TECHNIQUE IS UNRELIABLE BEYOND THIS POINT DUE TO COMPOUNDING OF ERRORS.

RELATIVE LIKELIHOOD

- HOW MANY TIMES IS EVENT A MORE LIKELY THAN EVENT B?

MULTIATTRIBUTE UTILITY THEORY

deNeufville, Chapter 20, Problem 1 – Solution



$$U_X(X) = \frac{1}{32}(X^2 - 4) \quad 2 \le X \le 6$$

$$U_Y(Y) = \frac{1}{6}\sqrt{Y + 36} \qquad -36 \le Y \le 0$$

a. $U(X,Y)$ **for all four corners of its domain are shown in the figure**

b. **Find by interpolation** $U(4,-36)$, $U(4,0)$, **and** $U(4,-20)$ **(See figure)**

$$U(4,-36) = 0.5\frac{1}{32}(4^2 - 4) = 0.5\frac{12}{32} = 0.5\frac{3}{8} = \frac{3}{16} = 0.1875$$

$$U(4,0) = 0.4(\frac{3}{8}) + 0.6 = 0.15 + 0.6 = 0.75$$

$$U(4,-20) = \frac{9}{16}(\frac{1}{6}\sqrt{-20+36}) + \frac{3}{16} = \frac{9}{16}(\frac{4}{6}) + \frac{3}{16} = \frac{6}{16} + \frac{3}{16} = \frac{9}{16} = 0.5625$$

c. **Write the general formula for** $U(X,Y)$

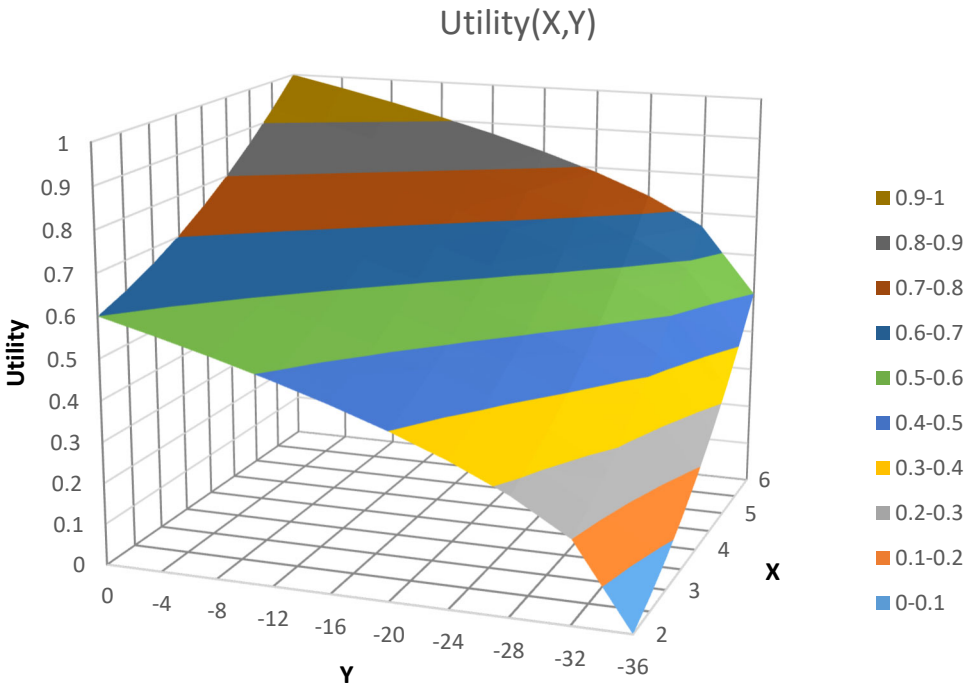
$$\begin{aligned} U(X,Y) &= k_X U_X(X) + k_Y U_Y(Y) + K k_X k_Y U_X(X) U_Y(Y) \\ &= k_X U_X(X) + k_Y U_Y(Y) + (1 - k_X - k_Y) U_X(X) U_Y(Y) \\ &= 0.5\left(\frac{1}{32}(X^2 - 4)\right) + 0.6\left(\frac{1}{6}\sqrt{Y+36}\right) + (1 - 0.5 - 0.6)\left(\frac{1}{32}(X^2 - 4)\right)\left(\frac{1}{6}\sqrt{Y+36}\right) \\ &= \frac{1}{64}(X^2 - 4) + 0.1\sqrt{Y+36} - \frac{1}{1920}(X^2 - 4)\sqrt{Y+36} \end{aligned}$$

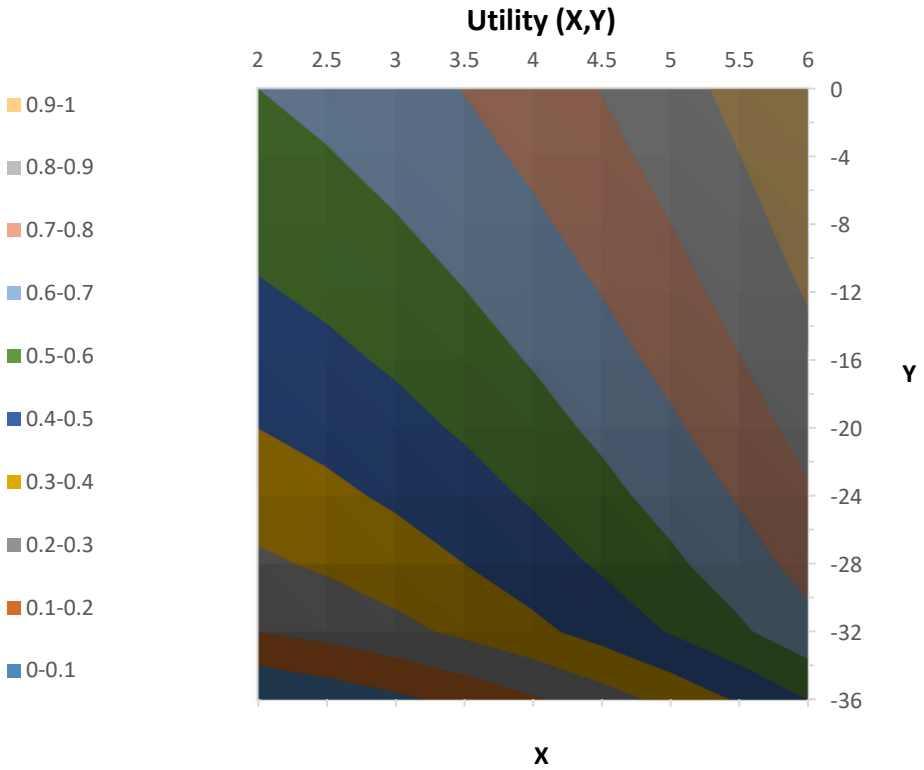
d. **Solve for the value of** K **in the formula**

$$K = \frac{1 - k_X - k_Y}{k_X k_Y} = \frac{1 - 0.5 - 0.6}{0.5 \times 0.6} = \frac{-0.1}{0.3} = -\frac{1}{3}$$

e. **Calculate** $U(4,-20)$ **by using the formula for** $U(X,Y)$

$$\begin{aligned} U(4,-20) &= \frac{1}{64}(4^2 - 4) + 0.1\sqrt{-20+36} - \frac{1}{1920}(4^2 - 4)\sqrt{-20+36} \\ &= \frac{12}{64} + 0.4 - \frac{48}{1920} = 0.1875 + 0.4 - 0.025 = 0.5625 = \frac{9}{16} \end{aligned}$$


Utility(X,Y)
Utility
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
-4
-8
-12
-16
-20
-24
-28
-32
-36
Y
2
3
4
5
6
X
0.9-1
0.8-0.9
0.7-0.8
0.6-0.7
0.5-0.6
0.4-0.5
0.3-0.4
0.2-0.3
0.1-0.2
0-0.1


Utility (X,Y)
2
2.5
3
3.5
4
4.5
5
5.5
6
X
0
-4
-8
-12
-16
-20
-24
-28
-32
-36
Y
0.9-1
0.8-0.9
0.7-0.8
0.6-0.7
0.5-0.6
0.4-0.5
0.3-0.4
0.2-0.3
0.1-0.2
0-0.1