
APPLIED SYSTEMS ANALYSIS

Engineering Planning and Technology Management

Richard de Neufville
*Massachusetts Institute
of Technology*

McGraw-Hill Publishing Company

New York St. Louis San Francisco Auckland Bogotá Caracas Hamburg
Lisbon London Madrid Mexico Milan Montreal New Delhi Oklahoma City
Paris San Juan São Paulo Singapore Sydney Tokyo Toronto

16.12 *P. O'Toole*

Paddy O'Toole makes outdoor furniture that sells in the summer (the Paddy O. Furniture line). Production starts in January, however, right after the holidays. That is when Paddy must decide on whether to hire 1, 2, or 3 assistants at \$14,000 each for the season. He could then produce Low (1200 units), Medium (1800 units), or High (2400 units). His product sells for \$30/unit on average in June. Units not sold in June must be sold at half price. Based on experience, he estimates $P(N)$, the probability that the demand for his furniture next June will be N , as:

N	:	1000	1500	2000	2500
$P(N)$:	0.2	0.3	0.4	0.1

How many assistants should Paddy hire to maximize expected profit?

CHAPTER 17

INFORMATION GATHERING

17.1 THE ISSUE

New information about a situation provides us with the opportunity to base our decisions on more current, better estimates of the consequences of any choice. We may expect that new information thus leads to better decisions. This logic suggests that any decision analysis should consider spending some effort gathering information.

Any decision problem has opportunities for gathering new information. It is collected through deliberate efforts to test the situation. For example

- An oil company may run detailed geological surveys before deciding whether to drill a test well.
- A manufacturer can build a prototype machine or even a factory before committing to full-scale production.
- A distributor may test market a new product before deciding if and how to sell it.
- A planner may use forecasting models to project future demand for a service.

In general, professionals are always engaged in some explicit forms of information gathering, so that they can improve their decisions.

Information can also be obtained simply by waiting. Indeed, the best way to determine what will happen N years hence is to wait until then and find out. If a decision must be taken earlier, one can at least wait awhile and see how some of the uncertainties about the future are resolved. There is thus a good reason to defer decisions.

In practical terms, the decisionmaker must decide how much information to obtain, if any. We must expect that additional information will be less and

less useful. The cost of obtaining it will at some point exceed its benefits. To determine this point, and thus to define the optimum amount of information to collect, it is necessary to define the value of the information. This chapter shows how the analyst can calculate the value of information, and use this result to improve decisions and, more generally, the strategy for a system.

17.2 CONCEPT

The *value of information* in decision analysis is: The increase in expected value to be obtained from a situation due to the information, without regard for the cost of obtaining it. This definition is usefully explained by example. Consider a risky situation, *L*. (This would also frequently be called a "lottery," a term defined precisely in Section 18.5.) The analyst can calculate the sequence of decisions which maximizes the expected value to be obtained from the situation. Call the value of the optimum strategy $EV^*(L)$. Consider next the same risky situation, altered by the influence of new information which has revised the estimated probabilities of various consequences. Call this new version *LI*, and the value of its optimum strategy $EV^*(LI)$. The definition is then:

$$\text{Value of Information} = EV^*(LI) - EV^*(L)$$

This definition is explicitly devised to fit the needs of the analyst. It first of all focuses directly on how the information affects the problem being analyzed. Second, it is defined before costs are considered so that it can be compared directly to these costs. This distinction permits the analyst to address such practical questions as

- Is this information worth its cost?
- What is the benefit-cost ratio of getting this new information?
- How much can I afford to pay for the information?


Because the value of any information as defined is intrinsic to a problem, whereas its cost may be negotiable, the separation between the value and cost in the definition provides the manager with a good basis for discussing what to pay for information.

Decision analysts generally refer to a *test* in discussing this problem. A test simply refers to any specific process of collecting information. It may be as simple as what we normally think of as a test, that is, a one-shot verification of a system, or it may be quite complicated. Note carefully that tests cannot, in general, conclusively demonstrate exactly how a system performs (see the discussion of "false positives" in Section 17.3).


Decision tree after test. To calculate the value of any information, it is necessary to understand the decision tree that results from a data-gathering activity or test. This is as follows. Any data-gathering activity will produce one of a set of *k* observations, or test results, TR_k , out of many possibilities. Each of these possibilities has some probability of occurring, P_k . The decision to collect information thus leads to the chance node illustrated (see box) as

The Decision Tree for the Test Results

To illustrate what happens after a test, consider the simple decision problem discussed in Section 16.5: Should you take your raincoat or not the next day you go to work or school? Suppose again that the risk is that it may, for simplicity, either rain or not, and that the value of the various possible outcomes (such as being caught in the rain without a raincoat, etc.) are the numbers indicated on the right, below. Your original decision tree is thus:




Suppose further that your data-gathering activity is to turn on the news to hear the weather forecast. For simplicity, suppose you might only hear either of two reports: rain predicted or not. From experience you could know the probability, P_k , of getting either report, TR_k , on a given day. Your decision tree for the test results would then be:



$$D \text{ --- collect information --- } C \text{ --- } P_k \text{ --- } TR_k$$


Each observation leads to a revision of the probabilities of events associated with the original risky situation, *L*. Call this revised risky situation, after test result TR_k , L_k . Graphically, as further explained in the next box,



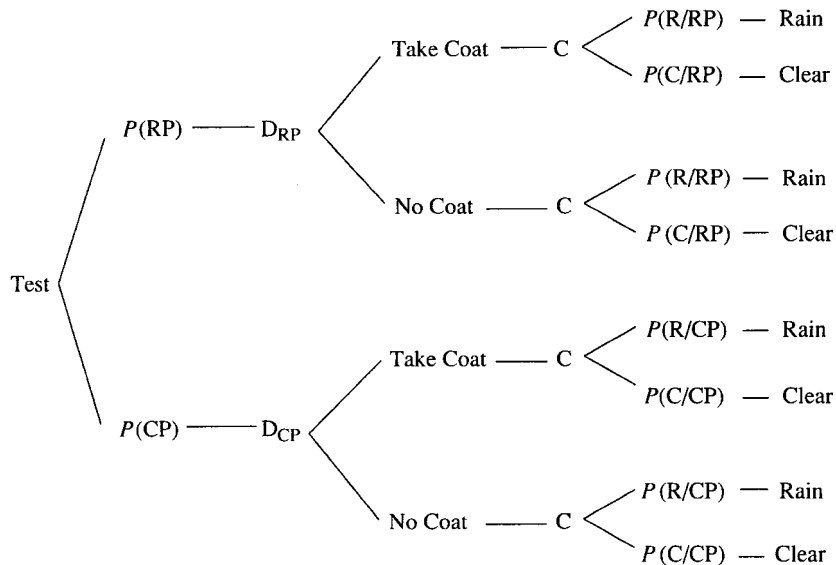
The risky situation represented by the decision to obtain information, *LI*, is the combination of the preceding two elements, the chance node for the test results,

Consequence of Test Results

Once you have the test result, either that rain is predicted or not, you have to revise your original decision tree. You of course have as many revisions as possible test results. For example, if the news report is "Rain Predicted," or RP, the subsequent decision tree is



The complete decision tree after the test is simple, conceptually. It is the aggregation of the above. Graphically it is a messy bush, which we only sketch. Using CP for "Clear Predicted," it is



and a second stage of decisions consisting of k different new versions of the original decision tree:

$$LI \equiv D \text{ --- collect information --- } C \text{ --- } P_k \text{ --- } L_k$$


The expected value of this situation is the expectation of the maximum value obtainable from each L_k and its probability of occurrence:

$$EV^*(LI) = \sum P_k EV^*(L_k)$$

Thus the value of information is:

$$\sum P_k EV^*(L_k) - EV^*(L)$$

Complete decision tree with test. The total problem facing the analyst considering whether to collect some particular kind of information, that is, to do a specific test, is the original decision tree augmented by a single test branch leading to as many revisions of this original decision as there are possible outcomes of the test. Graphically, it is



The value of information is the increase in value of the decision due to the addition of the test option. It is simply the difference between the expected value of the best sequence of decisions associated with the test compared to that associated with the original decision tree.

Events versus outcomes. It is important to notice that the result of any information-gathering activity or test is some knowledge about events E_j which are different from the outcomes O_{ij} that result from any decision. For our simple raincoat problem, for example, the information concerns whether it will rain or not; these are the two possible events. The outcome O_{ij} are the results that occur after making a decision and being subject to a chance event: for instance, deciding not to take a raincoat and being in the rain (event E_j) and getting wet (outcome O_{ij} , equaling -10 in the example).

Further examples of the differences between the events observed and the outcomes are

- An oil company might look for salt domes (event) to help them improve their forecast of oil reserves (outcome).
- A manufacturer would observe the productivity or sales (event) of a prototype, to make more realistic estimates of the profitability of full-scale production (outcome).

- A distributor would use the results of a test market (event) for a product in a particular city to estimate the market share that could be obtained from national sales (outcome).
- A planner would use a prediction made by a forecasting effort (event) to recalibrate the future loads on a system (outcome).

17.3 SAMPLE AND PERFECT INFORMATION

Two concepts of information must be considered: sample and perfect. Sample information is what you actually get from a test; its value is complicated to calculate. Perfect information is a totally hypothetical idea; its value is both easy to estimate and provides an upper bound to the value of information. Perfect information is thus a really useful artifact. Each type is discussed in detail below.

Sample information. Sample information is a set of actual observations as obtained in practice. It is inevitably a sample of the entire set of observations that might have been obtained if we had the resources to collect data indefinitely. It is necessarily incomplete.

Real, sample information also contains errors. Generally, these fall into two categories, either false positives or false negatives. A *false positive* is an erroneous report that an effect exists when it does not in fact. A “false alarm” for a fire is an example. False positives also occur routinely in medical examinations (see box). Conversely, a *false negative* is the false report that an effect does not exist when it actually does. An inspection that fails to detect a malfunction and thus reports that a system is in good order is an example of a false negative.

Because our instruments, our means for gathering data, are imperfect, we always have the possibility of both false positives and negatives. The probability of getting a specific test result, for example that an effect exists, is thus different from its actual probability of existence: $P(TR_k) \neq P(E_k)$ for sample information.

To calculate the probability of any test result, we must thus estimate the total probability of all the ways it may occur. These consist both of correct reports of an event when it occurs—a fraction of the whole since sometimes the event is not

Medical False Positives and Negatives

In many parts of the United States, teachers have been traditionally required by law to pass periodic tests for Tuberculosis (TB). This is a public health measure designed to protect children from contagion. TB or not TB? That is the question.

Screening for TB is routinely done by chest X-rays. Unfortunately, the dark spots on the lungs associated with TB can also be caused by other minor or noncommunicable diseases. Technicians may also of course either misread the photographs or transcribe the results erroneously. False positives are thus common with these X-ray examinations.

reported to exist when it really does, a false negative—and of the false positives. The reliability of the test, the degree to which it generates false positives and negatives, is a key element of this calculation.

The general formula is

$$P(TR_k) = \sum_i P(TR_k/E_i) P(E_i)$$

where (TR_k/E_k) is the correct report of an effect E_k when it occurs, and (TR_k/E_i) , $i \neq k$ is the false positive report of E_k when actually E_i exists. Note that the summations of true and false positives are weighted by both the reliability of the test, $P(TR_k/E_i)$ and the frequency of occurrence of each case, $P(E_i)$. (See following box for example calculation.)

Probability of Test Result

Suppose that the X-ray examination for TB in some lab has a 90% chance of being correct. That is, 9 times out of 10 TB will be reported (TBR) when the person has it, and conversely that a clean report (CR) will be given when the person is healthy:

$$P(TBR/TB) = P(CR/C) = 0.9$$

There is then a 10% chance of false positives and negatives:

$$P(TBR/C) = P(CR/TB) = 0.1$$

Note that this is a simple case. In general the probability of a test result, the reliability of our measuring device, may depend on what is to be found. For example, the X-ray technician may be more likely to miss TB spots than to exaggerate them.

Suppose additionally that TB is a rare disease among the population being tested, say that only 1% of the group has it. The rest of the people are “clean,” thus:

$$P(TB) = 0.01 \quad P(C) = 0.99$$

In this situation, the chance that anyone in the group will be reported as having TB is:

$$\begin{aligned} P(TBR) &= P(TBR/TB)P(TB) + P(TBR/C)P(C) \\ &= (0.9)(0.01) + (0.1)(0.99) \\ &= 0.009 + 0.099 = 0.108 \end{aligned}$$

Note that in this particular case the incidence of false positives, $P(TBR/C) P(C) = 0.099$, far outweighs the incidence of true positives, $P(TBR/TB) P(TB) = 0.009$. This is because the error rate applies to the bulk of the population, which is so much greater than the small number of individuals with the TB condition.

- A distributor would use the results of a test market (event) for a product in a particular city to estimate the market share that could be obtained from national sales (outcome).
- A planner would use a prediction made by a forecasting effort (event) to recalibrate the future loads on a system (outcome).

17.3 SAMPLE AND PERFECT INFORMATION

Two concepts of information must be considered: sample and perfect. Sample information is what you actually get from a test; its value is complicated to calculate. Perfect information is a totally hypothetical idea; its value is both easy to estimate and provides an upper bound to the value of information. Perfect information is thus a really useful artifact. Each type is discussed in detail below.

Sample information. Sample information is a set of actual observations as obtained in practice. It is inevitably a sample of the entire set of observations that might have been obtained if we had the resources to collect data indefinitely. It is necessarily incomplete.

Real, sample information also contains errors. Generally, these fall into two categories, either false positives or false negatives. A *false positive* is an erroneous report that an effect exists when it does not in fact. A “false alarm” for a fire is an example. False positives also occur routinely in medical examinations (see box). Conversely, a *false negative* is the false report that an effect does not exist when it actually does. An inspection that fails to detect a malfunction and thus reports that a system is in good order is an example of a false negative.

Because our instruments, our means for gathering data, are imperfect, we always have the possibility of both false positives and negatives. The probability of getting a specific test result, for example that an effect exists, is thus different from its actual probability of existence: $P(TR_k) \neq P(E_k)$ for sample information.

To calculate the probability of any test result, we must thus estimate the total probability of all the ways it may occur. These consist both of correct reports of an event when it occurs—a fraction of the whole since sometimes the event is not

Medical False Positives and Negatives

In many parts of the United States, teachers have been traditionally required by law to pass periodic tests for Tuberculosis (TB). This is a public health measure designed to protect children from contagion. TB or not TB? That is the question.

Screening for TB is routinely done by chest X-rays. Unfortunately, the dark spots on the lungs associated with TB can also be caused by other minor or noncommunicable diseases. Technicians may also of course either misread the photographs or transcribe the results erroneously. False positives are thus common with these X-ray examinations.

reported to exist when it really does, a false negative—and of the false positives. The reliability of the test, the degree to which it generates false positives and negatives, is a key element of this calculation.

The general formula is

$$P(TR_k) = \sum_i P(TR_k/E_i) P(E_i)$$

where (TR_k/E_k) is the correct report of an effect E_k when it occurs, and (TR_k/E_i) , $i \neq k$ is the false positive report of E_k when actually E_i exists. Note that the summations of true and false positives are weighted by both the reliability of the test, $P(TR_k/E_i)$ and the frequency of occurrence of each case, $P(E_i)$. (See following box for example calculation.)

Probability of Test Result

Suppose that the X-ray examination for TB in some lab has a 90% chance of being correct. That is, 9 times out of 10 TB will be reported (TBR) when the person has it, and conversely that a clean report (CR) will be given when the person is healthy:

$$P(TBR/TB) = P(CR/C) = 0.9$$

There is then a 10% chance of false positives and negatives:

$$P(TBR/C) = P(CR/TB) = 0.1$$

Note that this is a simple case. In general the probability of a test result, the reliability of our measuring device, may depend on what is to be found. For example, the X-ray technician may be more likely to miss TB spots than to exaggerate them.

Suppose additionally that TB is a rare disease among the population being tested, say that only 1% of the group has it. The rest of the people are “clean,” thus:

$$P(TB) = 0.01 \quad P(C) = 0.99$$

In this situation, the chance that anyone in the group will be reported as having TB is:

$$\begin{aligned} P(TBR) &= P(TBR/TB)P(TB) + P(TBR/C)P(C) \\ &= (0.9)(0.01) + (0.1)(0.99) \\ &= 0.009 + 0.099 = 0.108 \end{aligned}$$

Note that in this particular case the incidence of false positives, $P(TBR/C) P(C) = 0.099$, far outweighs the incidence of true positives, $P(TBR/TB) P(TB) = 0.009$. This is because the error rate applies to the bulk of the population, which is so much greater than the small number of individuals with the TB condition.

One point needs to be carefully noted in connection with the results of real tests: no test can prove that any theory is correct. When a test result is compatible with theory it merely increases the probability that the theory may be correct. This is because no real test, with its imperfections, can exclude false positives and negatives. The change in credibility can be calculated by applying Bayes' Theorem (see Section 15.3 and following box). This point needs to be stressed because it confuses most people, who naturally—but wrongly—believe that a positive result confirms theory.

This confusion over the interpretation of test results is especially significant when we are trying to detect relatively rare conditions (such as AIDS) with imperfect measurements. What happens is that positive test results may easily be mostly due to the false positives from the bulk of the population that does not have the condition.

Positive tests on rare events should thus be repeated. This is standard procedure in medicine: positive results on TB, for example, are followed up with more precise (and more costly) tests. Similar procedures should be applied to all information gathering on important rare events.

Perfect information. This is an artificial concept that is most useful in practice. It represents the test results that would come from a hypothetical process that identifies exactly which event is true or what will happen. For example, perfect information about the incidence of TB would specify exactly which individuals had it and which did not; no false positives or negatives. Perfect information would give us exact knowledge about the environment and the consequences of decisions we might make.

The concept of perfect information provides an upper bound on the value

Misinterpretation of Test Results

A positive report from the X-ray test, that someone has TB, does not demonstrate that the person is actually affected, contrary to one's immediate impression. We can calculate the new probability of having the disease, given this new information, by Bayes' Theorem. Using our previous data,

$$\begin{aligned} P(\text{TB}/\text{TBR}) &= P(\text{TB})[P(\text{TBR}/\text{TB})/P(\text{TBR})] \\ &= 0.01[0.9/0.108] \\ &\sim 0.09 \ll 1.00 \end{aligned}$$

In this case, as so often with rare phenomena, the number of false positives dominates the test results (see previous box). Even though the test is relatively accurate, and revises the probability by an order of magnitude, it is still far from proving the condition.

of information from any test. Since perfect information—were it to exist—would be the best possible set of observations, it must have the highest value.

The value of perfect information is easy to calculate, far easier than the value of real, sample information. The next section describes this in detail. This feature, together with the fact that perfect information provides an upper bound, means that perfect information provides a practical basis for initial estimates of the value of any information-gathering process.

17.4 VALUE OF INFORMATION

We can calculate the value of both perfect and sample information. We consider the value of perfect information first because it is easier to calculate, more widely used, and often all that is necessary.

Expected value of perfect information (EVPI). The value of perfect information is easy to calculate. Perfect information both radically simplifies the decision tree after the test, and makes it trivial to assign the probabilities throughout. Neither revised probabilities of events nor the probabilities of the test results have to be calculated.

First, every test result, TR_k , from the perfect test will tell us exactly what will happen subsequently, and its associated outcome, O_{ik} , will have probability one in the revised decision tree following the test result, L_k : $P(\text{O}_{ik}) = 1.0$ and all other outcomes, (non- O_{ik}), will have probability zero, $P(\text{non-}\text{O}_{ik}) = 0$.

The optimal decision for each L_k , (D/TR_k), should then be obvious: one simply takes the choice with the highest expected value, which is then $EV^*(L_k)$.

Second, the probability of making any observation TR_k is obvious when we are dealing with perfect information. The estimated probability of making a perfect observation that a test result indicating that event E_k will occur is identical to our prior probability of E_k itself: a perfect test would by definition predict that E_k would happen exactly as many times as it does.

The value of perfect information follows directly from the two simplifications described above. The expected value after the information-gathering process is:

$$EV^*(\text{LI}) = \sum P(E_k)EV^*(\text{D}/\text{TR}_k)$$

The expected value of perfect information is then:

$$EVPI = \sum P(E_k)EV^*(\text{D}/\text{TR}_k) - EV^*(L)$$

For most practical problems, where we already have estimates of the probabilities of events, this calculation can practically be done by inspection (see Section 17.5).

The EVPI upper bound on the value of information is often all that is required in practical situations. It constitutes a minimum test any proposed data-gathering project should meet, and generally eliminates many alternatives from

consideration. Indeed, if the project—a prototype plant, a geological exploration, or whatever—costs more than the maximum conceivable benefits, then it is clearly not worth doing. In practice, it is reasonable to impose more stringent standards, since no real information is likely to be near perfection. As a rule of thumb, providing a reasonable maximum to the value of sample information, we may usefully require the cost of any data-gathering project to be less than half the expected value of perfect information.

Expected value of sample information (EVSI). The expected value of sample information is calculated as outlined above—without any simplifications! Even for trivial textbook situations it is hard work.

The calculations involve several complications. First, revised probabilities of the outcomes, $P(O_{ik}/TR_k)$, must be calculated for each possible set of observations TR_k : k modifications of Bayes' Theorem. Then, each new risky situation L_k must be analyzed to obtain the best sequence of decisions for L_k and thus $EV^*(L_k)$: k modifications of the original decision analysis for L . Next, the probabilities of making any observation TR_k must be calculated:

$$P_k = P(TR_k) = \sum_j P(TR_k/E_j) P(E_j)$$

Finally, all this must be put together to obtain $EV^*(LI)$. The expected value of sample information is then

$$EVSI = \sum_k P_k EV^*(L_k) - EV^*(L)$$

The example in Section 17.5 illustrates what is involved. Even when computers are available, the effort is sufficiently complex to be avoided if possible. This underlines the usefulness in practice of the concept of perfect information.

17.5 EXAMPLE CALCULATIONS

This section shows how to calculate the expected values of perfect and sample information for the example problem given in Section 17.2. The information, in this case, consists of the weather forecast heard on the news, which can either be "Rain Predicted" or "Clear Predicted."

Expected value of perfect information (EVPI). Perfect information would mean—in this hypothetical case—that the weather forecast would predict exactly what will happen: if it says "Rain Predicted," it will rain.

Perfect information simplifies the decision tree after the test in two ways. First, it makes it possible to know the probability of getting any test results. In our case, our best estimate of the probability of getting a "Rain Predicted" forecast exactly equals our prior estimate of the probability of rain:


$$P_{RP} = P(\text{Rain}) = 0.4$$

The hypothetical perfect forecast would say "Rain Predicted" every time it would rain and only those times. Similarly for the probability of other possible test results. In that case,

$$P_{CP} = P(\text{Clear}) = 0.6$$

The second simplification provided by perfect information concerns the best decision after you hear the test results. These would tell you what will happen, and it then should be obvious what choice to make. In our case, if you heard "Rain Predicted" from the perfect forecast, you would know it will rain and that taking your raincoat was the best you could do at that point. Likewise, if the perfect forecast said "Clear Predicted" you would know not to take your raincoat.


The decision tree after the perfect forecast is thus very simple. In our case it is



The expected value of having the test is thus

$$EV^*(LI) = 0.4(5) + 0.6(4) = 4.4$$

To find the expected value of perfect information, we must compare the expected value after the test with the expected value we would have had without the test. Our decision tree without the test was, from 17.2,



The expected values of the alternatives are, as shown in Section 16.5,

$$EV(\text{Take Coat}) = 0.8 \quad EV(\text{No Coat}) = -1.6$$

Taking the raincoat is the better decision, with the highest expected value, which is the expected value of the situation without the test:

$$EV^*(L) = 0.8$$

The expected value of perfect information is, by definition, the difference between the value with and without the test:

$$EVPI = 4.4 - 0.8 = 3.6$$

Expected value of sample information (EVSI). Calculating the EVSI requires 5 steps:

1. Obtain the probabilities of the test results, P_k .
2. Revise the prior probabilities of the outcomes, based on the test results TR_k , to get the posterior probabilities $P(O_{ij}/TR_k)$. Note this requires you to apply Bayes' Theorem for each $P(O_{ij})$ for each TR_k .
3. Identify the best decision you could make after each test result, D_k^* . Note that this requires you to repeat k times the original analysis of the decision without test information, each time with different probabilities.
4. Calculate the expected value after the test, as the expected value of each optimal decision after a test result times the probability of that test result:

$$EV(\text{after test}) = \sum_k P_k EV^*(L_k)$$

5. Finally, find the expected value of sample information as the difference between the above and the expected value without the test.

The whole process is a big job—especially as compared to the ease of calculating the expected value of perfect information.

To obtain the probabilities of the test results, we need to know something about the reliability of the test. This we can usually obtain from an analysis of the past performance of the test. Suppose for simplicity that the weather forecast has a 70% chance of being correct:

$$P(RP/R) = P(CP/C) = 0.7$$

and conversely, then, a 30% chance of predicting something different:

$$P(RP/C) = P(CP/R) = 0.3$$

The possibility of each test result is calculated directly from the information on the reliability of the test, as described in Section 17.3. Thus for our case,

$$\begin{aligned} P(RP) &= P(RP/R)P(R) + P(RP/C)P(C) \\ &= (0.7)(0.4) + (0.3)(0.6) = 0.46 \end{aligned}$$

and in this simple case of two possibilities,

$$P(CP) = 1 - P(RP) = 0.54$$

Second, we must revise our prior probabilities of the outcomes depending on each of the different test results. In our case this is twice. If the forecast says "Rain Predicted,"

$$P(R/RP) = P(R) \left[\frac{P(RP/R)}{P(RP)} \right] = 0.4 \left(\frac{0.7}{0.46} \right) = 0.61$$

$$P(C/RP) = 1 - 0.61 = 0.39$$

If, however, the forecast says "Clear Predicted,"

$$P(C/CP) = 0.6(0.7/0.54) = 0.78$$

$$P(R/CP) = 1 - 0.78 = 0.22$$

Third, we must identify the best decision we could make given the test result and the consequently revised probabilities. Thus if we have "Rain Predicted," we can calculate the expected value of each alternative as

$$EV(\text{Take Coat})_{RP} = (0.61)5 + 0.39(-2) = 2.27$$

$$EV(\text{No Coat})_{RP} = (0.61)(-10) + 0.39(4) = -4.54$$

The best choice for "Rain Predicted" is to take the raincoat, and the value of that optimal decision is:

$$EV^*(L_{RP}) = 2.27$$

If, however, we have "Clear Predicted" we have

$$EV(\text{Take Coat})_{CP} = (0.22)5 + (0.78)(-2) = -0.46$$

$$EV(\text{No Coat})_{CP} = (0.22)(-10) + (0.78)(4) = 0.92$$

The best choice here is not to take the coat, so

$$EV^*(L_{CP}) = 0.92$$

Now, finally, we can calculate the expected value after the test as the sum of the probability of each test result times the value of the best decision after that result:

$$EV^*(LI) = \sum P_k EV^*(L_k) = 0.46(2.27) + 0.54(0.92) = 1.54$$

The expected value of sample information can now at last be found:

$$\begin{aligned} EVSI &= EV^*(LI) - EV^*(L) \\ &= 1.54 - 0.8 = 0.74 \end{aligned}$$

This result and process should be carefully compared to those of the expected value of perfect information. First, perfect information provides an upper bound. Furthermore, as a rough rule of thumb indicated earlier, perfect information is generally at least twice as valuable as what you might actually get in a real test. In our case we have:

$$EVPI = 3.6 > EVSI = 0.74$$

The process for getting the expected value of perfect information is also very, very much simpler—even for the trivial case just presented. This reality

underlines the advantage of calculating the value of perfect information and using it, where possible, at least as a first means for evaluating the desirability of any test or information-gathering procedure.

17.6 APPLICATION

Developers in the San Francisco area planned to build on about 350 acres or 1.4 km² of land reclaimed from the Bay. Their problem was to decide among some combination of three alternatives:

1. Immediate construction, without waiting for the soil to consolidate
2. Preloading of the site, to accelerate the consolidation
3. Waiting until consolidation occurred naturally

The risk they faced was that the soil might not have consolidated sufficiently, and might then settle excessively under the buildings, causing cracks and other damage. This risk depended on the compressibility of the soil, which may be considered the chance event for this decision analysis.

New information relevant to the analysis would come from samples of soil obtained from borings. The general question is whether it is worthwhile to carry out a test program (however boring). Specifically, it is appropriate to ask how many borings to make, in other words, which test program maximizes the value of information.

The value of the information naturally depends on the quality of the information already available, that is, the prior probability distribution. That in turn depends on the experts who provide it. In this case, three types were available: Senior professionals with an average of 15 years experience, junior engineers with 3 years on the job, and current graduate students.

Colleagues at Stanford University calculated the value of information for each type of expert and for tests involving different numbers of soil samples. They used continuous probability distributions such as those referred to in the box in Section 15.4. Once they had the value of information, they calculated the net gain associated with any program by subtracting off the cost of the borings. Their results are as shown in Figure 17.1.

The value of information in this case was minimal for experienced local engineers. It was significantly less than the cost of the test programs, which were thus not worthwhile. As frequently happens, in this and other fields, two complementary factors were at work: the experts understood the situation and the new information did not noticeably change the prior estimates. These factors combine to reduce the value of information.

Contrarily, the value of information increases as the priors degrade. It is worth more for junior than senior engineers, and the most for students. Because it is worth more, more tests are optimal for the less experienced. The optimal sample size for the three groups goes from 0 to around 5 to 15.

The fact that new information about soils has little or no value to soils experts in an area is well known to practicing professionals. Yet we can also

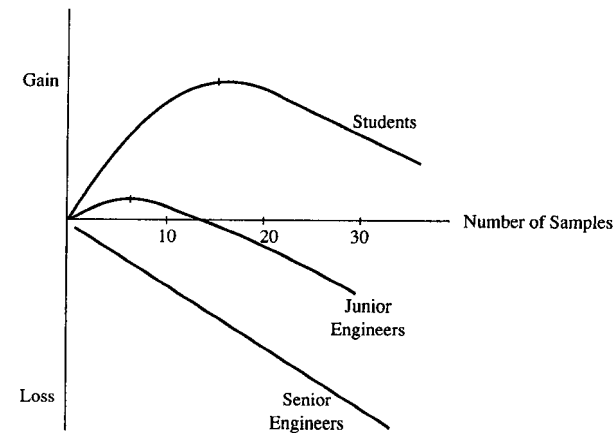


FIGURE 17.1
Expected net gain (value of information less cost of samples) from soil tests in example.

observe that they continue to collect it routinely. Why is this? They do so to protect themselves from lawsuits and claims of negligence. Even though the information may have no influence on the design, it may be necessary documentation for display in court. This kind of wasteful effort is an unfortunate, but common, result of a litigious society. It is known as “defensive” engineering, similar to “defensive” medicine of doctors.

17.7 CONSEQUENCES FOR STRATEGY

Whenever substantial risk is involved, the best approach to systems design involves a strategy, as Chapter 16 indicates. This consists of an optimal choice for the first phase of implementation, which incorporates sufficient flexibility so that the designer can alter the development to correspond to how the future unfolds.

Information gathering is important both to the choice of the first phase and to the implementation of the strategy. The collection of relevant additional information, that would clarify major uncertainties about a system, should often in fact be the first phase.

The first step in the design of a system should often be a major “testing” program. The information will have value for the optimal design, and should be obtained. As indicated in the beginning of this chapter, these tests may involve all kinds of surveys, investigations, trials, or prototypes. They each have the same common element: they establish progress in the development of a system with a minimum commitment.

The advantage of this approach can be appreciated by imagining how differently the U.S. utilities would have fared if they had consciously planned to

develop nuclear power using information gathering as a major component of their plans. They would have soon appreciated that nuclear plants were far more expensive to build and operate than they had thought, and could have altered their plans accordingly. Having failed to do so, several major U.S. utilities faced bankruptcy.

In implementing a strategy, the practical question is: When should one proceed with the next phase of development? Here again, a conscious program of information gathering, directed toward revising the assumptions and estimates of risks on which the initial phase was built, is the key to timely, appropriate decisions.

The analysis of the value of the information will, in each phase of the system development, indicate how extensive a test program should be. It is thus an integral part of good systems design in risky situations.

REFERENCES

- Behn, R. D., and J. W. Vaupel, (1982). "Compound Decision Dilemmas," Chapter 10, *Quick Analysis for Busy Decision-Makers*, Basic Books, New York.
- Folayen, J., K. Hoeg, and J. R. Benjamin, (1970). "Decision Theory Applied to Settlement Predictions," *Journal of Soil Mechanics and Foundations*, American Society of Civil Engineers, Vol. 96, pp. 1127-1141.
- de Neufville, R., and R. Keeney, (1972, 1974). "Use of Decision Analysis in Airport Development for Mexico City," Chapter 23, *Analysis of Public Systems*, A. W. Drake, R. Keeney, and P. Morse, eds., MIT Press, Cambridge, MA; and Chapter 24, *Systems Planning and Design*, R. de Neufville and D. H. Marks, eds., Prentice-Hall, Englewood Cliffs, NJ.
- Raiffa, H., (1970). "More About the Economics of Sampling," Chapter 7, *Introductory Lectures on Choices Under Uncertainty*, Addison-Wesley, Reading, MA.

PROBLEMS

17.1. Money Bags, Take 3

Welcome back (see Problem 16.2) to the "Money Bags" TV show. Determine the value of information about the contents of the wallet.

- The EVPI.
- The EVSI. Is it greater or less than the amount Monty wants you to pay for a look at the bills?

17.2. Assembly Robot (Revisited)

See Problem 16.3. Examine the value of information of the laser test.

- What is its EVSI? What is the longest time the test could take and still be worth doing?
- What is the EVPI of a perfect test?

17.3. Mars Probe Simulation

The scientists described in Problem 16.4 have the option, at a cost of \$1 million, of placing a remote sensing device on the probe to collect data about the Martian surface while the probe is still in orbit. The data could be used to simulate perfectly

a method X experiment, which could then be used in choosing the method for the actual experiment.

- Calculate the probability that theory A is correct if (1) the simulation is a success; (2) the simulation is a failure.
- What is the probability that the experiment will be a success given that the simulation fails and the scientists choose method Y?
- Draw the decision branch for deciding to run the simulation. Calculate the remaining probabilities and label the branches.
- Should the scientists pay for the remote sensor if their objective is to maximize expected monetary value?
- What is the maximum they should pay for this device?
- What is the EVPI about the correct theory about the Martian surface?

17.4. Software Development, Part 2

As a member of the board of directors of Celestial Software, (see Problem 16.5) you are reviewing their proposed decision to buy a software license.

- What is the EVPI about the success of the computers?

You can hire an industry analyst to predict which PC will be a hit. By comparing the analyst's record for choosing hits from machines similar to A and B, you have determined that if A actually becomes a hit, there is a 60% chance that the analyst will have predicted this. Also, if B turns out to be the hit, there is an 80% chance that the analyst will have said so. The analyst's services may be purchased for \$100,000.

- Calculate the probability that the analyst will predict A to be the hit.
- Calculate the probability that A will be the hit given that the analyst picks A; and that B will be the hit given that the analyst chooses B.
- Draw the decision branch for deciding to hire the analyst, and label all probabilities and outcomes.
- Should you hire the analyst?
- What is the maximum you would pay the analyst?

17.5. PY-RIC, again

Vic's lab (see Problem 16.10) suggests that the booster be checked by a procedure that predicts correctly if today's launch is successful or not 80% of the time. $P(\text{predicts success}|\text{success}) = P(\text{predicts failure}|\text{failure}) = 0.80$.

- Calculate the probability of a successful launch today if the test "predicts success;" if it "predicts failure." Use Bayes' Theorem.
- Calculate the probability of a successful launch today if two independent tests contradict each other: one "predicts success," the other "predicts failure." Vic must now decide whether to test.
- What is the upper bound on the economic value of a test?
- Since the countdown is progressing, Vic does not have time to calculate the value of the proposed test. What is the maximum he might reasonably commit to it?

17.6. Tax Shelter

Ten years from now, having made your first \$10 million in the engineering specialty of your choice, you must decide how to dispose of a tract of land originally purchased to avoid taxes. Your options are

- Spend \$3M to develop resort (option R)
- Spend \$1M to open a strip mine for coal (option M)
- Sell the property for \$0.1M (option S)

The outcomes from options R and M depend on whether a coal seam runs through your property. A geologist friend estimates a 75% probability that there is a coal deposit. He believes that if coal is present it will be on your neighbor's land as well.

If there is no coal, the resort would return \$12M in revenue, the mine would make nothing. If there *is* coal, your neighbor will eventually find out about it and start mining, spoiling the view and forcing you to convert the resort to a school worth just \$1M in revenue, but the mine would have made \$1M profit.

(a) Structure the decision tree. Select your best strategy.

One of your friends recommends a geological testing service whose test results will indicate either HIGH probability of finding coal, or LOW. Coal-bearing deposits will result in a HIGH test report 76% of the time; noncoal-bearing strata yield HIGH results in only 12% of the cases.

(b) What is the most you would pay to test for coal?

(c) Construct the decision tree for the test, and find the expected value of the test. Is it worth \$1M?

17.7. *Lost Records*

You have the opportunity to purchase one of two plots of land for \$4M. Their value depends on what you can build on them. Last year engineers took a set of 10 borings from each plot. In one set, 3 borings showed a hard rock base and 7 showed compressed soil, indicating that this plot will support a highrise building, making it worth \$6.5M. In the other set 5 borings showed loose soil and the other 5 compressed soil, indicating that this plot will only support houses, making it worth \$3M. The difficulty is that someone has lost the record of which set of samples goes with which plot.

The owner offers you the opportunity to obtain—at a cost of \$0.5M—one more boring from either plot, provided you commit to buy at least one plot that you can choose afterwards. From experience, you estimate the probability of getting a boring of a specific type is directly proportional to the number in the original set of samples. For example, the probability of getting a boring showing compressed soil is 0.7 from the \$6.5M plot.

(a) What is the expected value of perfect information on the plots?

(b) What is the most you would pay for a new boring?

(c) What is your best strategy?

17.8. *Bill Durr*

Bill Durr, president of a construction firm, proposes to set up a subsidiary. He estimates that he has an equal probability of making an \$800,000 profit or a \$400,000 loss. His friend, Pat Ronach, says to him: "By interviewing people at City Hall, I can determine whether or not the public works program will be expanding, staying the same, or cutting back. This should help you decide if you want to start the new business. For you, I'll only charge \$40,000 for the service."

- (a) What is the most he should conceivably pay for information about whether he would make a profit or loss?
- (b) Bill supposes that he has an 80% chance of making a profit if the city expands public works, the same chance of losing if it contracts, and an even chance of winning or losing if the level stays the same. He also estimates each situation is equally likely. What should he do, if he has no help from Pat?
- (c) Bill hires Pat and gets told: "There's a 40% chance of expansion—if that doesn't happen, it's equally probable that the public works program will contract or stay the same. Please pay \$40,000." Is Bill ahead or behind?

17.9. *Sonny and Ex Reyes*

Sonny's former wife, Ex, comes along and says she can do research that will investigate the probability of success for the new method of producing PV panels (see Problem 16.9). Ex wants \$1 million for her research.

(a) Draw the decision after the test.

(b) Should Sonny pay Ex Reyes's price? Explain.

17.10. *Penalty Clause*

Scientific Instruments Inc. (SII) has received an order with the stipulations that if the system arrives late, SII will pay \$1000 for each late day; if it is defective, the client will not pay for it.

SII will make \$40,000 if the system is delivered on time, but will lose \$17,000 if it is defective. SII can easily deliver the system on time. However it may, with $P = 0.5$, be defective. This deficiency can be removed by an adjustment that costs \$2000, and would delay delivery 10 days. SII can also test if the component needs adjustment. However, this test is fallible and costs something. Bob Ability estimates that there is a 0.6 probability that the test will be favorable (the component does not seem to need adjustment). He also thinks that if the test is favorable, there is one chance in four that the component does need adjustment; if the test is unfavorable, this probability increases to seven chances in eight.

(a) Draw the decision tree for SII, indicating probabilities.

(b) What is the highest amount SII should pay for a perfect test on the component?

(c) What is the highest amount SII should pay for a test?

(d) What strategy would you recommend SII follow?

17.11. *Mine Shaft*

A mine shaft has collapsed. Twenty miners were in the shaft, all in a single elevator. All may still be alive; or all may have died, either crushed by falling rock or poisoned by gas.

You have to decide whether to send in a rescue team. Excavating the rubble fast enough to help the victims is dangerous—an average of 5 rescuers would die. In addition, you believe that there is only one chance in five that the miners are alive.

(a) Structure this problem as a decision tree.

(b) What should you do to minimize expected deaths?

It would be less dangerous to send a few volunteers down to test the air. If they find poison gas, the miners are dead; if not, there is a 50% chance that the miners are alive.

- (c) What is the absolute maximum number of rescuers you should sacrifice to find out if the miners are alive?
- (d) What is the expected value of the gas test? Should you do it if it would cost one life?

17.12. Bridge Design

You are designing a bridge for an area in which experts believe that there is a 50-50 chance that a single major earthquake will hit the area in the next 30 years. The choice is between a conventional design, sure to collapse under the stress of a major earthquake, and a new design.

A scale model of the new design has collapsed in several laboratory tests simulating a major earthquake. The engineers attribute the weakness to either the steel or the concrete. Both theories are equally plausible. They agree, also, that the bridge's probability of failure due to a major earthquake is 20% if the fault is in the steel, and 60% if in the concrete.

A conventional bridge costs \$3M, and one with the new design \$3.5M. The cost of rebuilding the bridge has a present value of \$2.0M. Early warning systems for earthquakes will prevent any loss of life due to collapsing of the bridge.

- (a) Structure the decision tree for deciding which design should be used. Label the branches and indicate the consequences of each path. Calculate the probabilities of each chance event occurring and place these on the appropriate branch.
- (b) Assuming the client wants to minimize expected cost, which design is preferred?

An engineer discovers a unique field experiment, costing \$0.02M, that simulates the effects of an earthquake on the new design exactly, and the model will either collapse or not.

- (c) Draw a tree for deciding whether they should conduct the new experiment. Label the branches.
- (d) If the model collapses, what is the probability that the steel theory is correct? What if there is no collapse?
- (e) Calculate the revised probability that the earthquake results in a collapse.
- (f) What is the expected value of the field experiment? Explain this result. Should this experiment be conducted?
- (g) Redo (h) assuming that the new bridge design costs \$3.6M rather than \$3.5M. Explain why the result is or is not compatible with the previous explanation.
- (h) What is the expected value of perfect information about the correct theory? Assume the cost of the new design is \$3.6M.

17.13. Oil Well

An oil explorer must decide whether to drill on a site before her option expires or to abandon her rights to the site. Drilling a well costs \$100,000. If oil is struck, the explorer will sell and make a \$500,000 profit.

The explorer can hire a geologist to take a seismic sounding of the site at a cost of \$10,000. Such soundings reveal, with certainty, whether the geophysical substructure of the land is of type A, type B, or type C, where each type has a different degree of favorability for the existence of oil. From past experience, the probability that a certain substructure will occur, given that oil exists, is

$$P(\text{type A/oil}) = 14/28$$

$$P(\text{type B/oil}) = 9/28$$

$$P(\text{type C/oil}) = 5/28$$

The *a priori* probabilities for substructures at this site are

$$P(\text{type A}) = 0.20$$

$$P(\text{type B}) = 0.30$$

$$P(\text{type C}) = 0.50$$

Given that $P(\text{oil}) = 0.28$,

- (a) Draw a fully labeled decision tree for deciding whether (1) the explorer should make the seismic tests and (2) should drill for oil or abandon the site following the decision to test or not. In taking the seismic test, the decision about drilling will be made knowing the substructure that exists.
- (b) Calculate the monetary outcomes (use the explorer's change in assets) of each path on the tree and place each at the end of appropriate path. Remember that taking the seismic soundings incurs extra costs.
- (c) Taking the seismic soundings improves knowledge about the existence of oil. Calculate the probabilities of striking oil in each of the three types of substructures (i.e., $P(\text{oil/type A})$, and so on).
- (d) Label each chance branch in the tree from (a) with its probability of occurring.
- (e) What strategy should the explorer follow? To pay for the seismic test? To drill for oil without taking the test? What decision should be made following the test?
- (f) What is the maximum amount the explorer should pay for the seismic soundings? (i.e., what is the expected value of sample information?)
- (g) What is the expected value of perfect information about the existence of oil?

17.14. Poker Game

Brette is in the middle of a poker game and the pot is now \$175. Brette must either pay \$25 into the pot to "see" Bart's hand, or fold and pay nothing. Brette has a fair hand and knows she can win if Bart is bluffing. If Bart is not bluffing, Brette knows she cannot win.

Brette believes that Bart is bluffing with probability of 0.2. Therefore, if Brette "sees" and wins, she will get \$200. If she folds, she will lose nothing above what she has already paid into the pot.

Brette makes poker decisions on the basis of expected value.

- (a) Draw a decision tree for Brette, showing decision nodes and branches, chance nodes and branches, the relevant probabilities, and the outcomes.
- (b) What should Brette do?
- (c) What is the maximum Brette would pay to know whether Bart is bluffing or not?

Brette knows that Bart has strange habits when drinking and playing cards simultaneously. She has noticed from past experiences that Bart often drinks more quickly when he is bluffing. Brette believes that the probability that Bart is bluffing, given that he is drinking quickly, is 0.9. She believes the probability that Bart is bluffing, given that he drinks slowly, is 1/20. Brette is debating whether to buy Bart a drink.

- (d) If Bart has a drink, what is the probability that he will drink it quickly?
 (e) What is the maximum Brette should pay to buy Bart's drink?

17.15. House for Sale

You must sell your house by the end of the month. You now have an offer of \$200,000, but you think that by the end of the month there is a 60% chance that you will get an offer of \$240,000 (and a 40% chance that you will not get any new bids). If you turn down the offer of \$200,000, and do not get any new bids, you have a standing offer of \$160,000

Alternatively, you can have a real estate broker sell your home for you. The charge would be \$12,000. For this you also get the broker's estimate of the sale price. You would then have to tell the \$200,000 bidder whether you accepted that bid or whether you were going to let the broker find a buyer. Before you pay, however, the broker tells you it is equally likely that he will predict that he will be able to find you a bid of \$240,000 or a bid of \$160,000. Experience shows that the broker is correct half the time.

Assume that your objective is to get as much money for the house as possible.

- (a) Structure your situation as a decision tree.
 (b) What is the best decision you can make without hiring the broker?
 (c) What is the expected value of perfect information? (Think of the broker's predictions as a "test" of the future.)
 (d) Calculate the expected value of sample information. Explain your result.

17.16. Procter N. Gamble

Procter, president of a food company, must decide whether to market a new breakfast drink which the R and D division has developed. A special meeting devoted to this topic yields the following information:

- The marketing vice-president has defined two possible outcomes for the success of this product; either the public will accept the product or it will not. She believes that the product will be accepted with probability 0.1.
- The cost engineers believe that if the product is marketed and accepted, the company will net \$100,000 yearly. If the product is rejected, however, the company will suffer a net loss of \$20,000 yearly. If Procter decides not to market the product, her company will neither accrue more cost nor make any profit on this product.
- Procter always makes decisions based on the expected value of the outcomes.

- (a) What is the best strategy in this case?
 (b) What is the value of a test which would reveal perfectly whether the public would accept or reject the breakfast drink?

Before she can implement the strategy found in part (a), the marketing VP informs Procter that she has developed a market test for the product. She proposes to test a sample population to determine its preference for the breakfast drink. Based on histories of similar tests, she estimates that the *sample population* would not accept the drink with a probability 0.2, when in fact the product would be accepted by the general public. She also believes that one time in forty the drink

would be accepted by the public when the test group rejects it. Finally, she informs Procter that the probability that the public accepts the drink when the test group accepts is 0.4.

- (c) Sketch the decision tree for the market test.
 (d) Calculate the probability that the test appears favorable.
 (e) What is the maximum Procter should pay for this test?

17.17. Marian Haste, Again

- (a) See Problem 16.11. What is the maximum Marian should pay for better information about the weather?
 (b) See Problem 15.8. What is May's advice worth to Marian?

17.18. P. O'Toole, Again

See Problem 16.12. What is the most Paddy should pay for a better forecast of June sales?