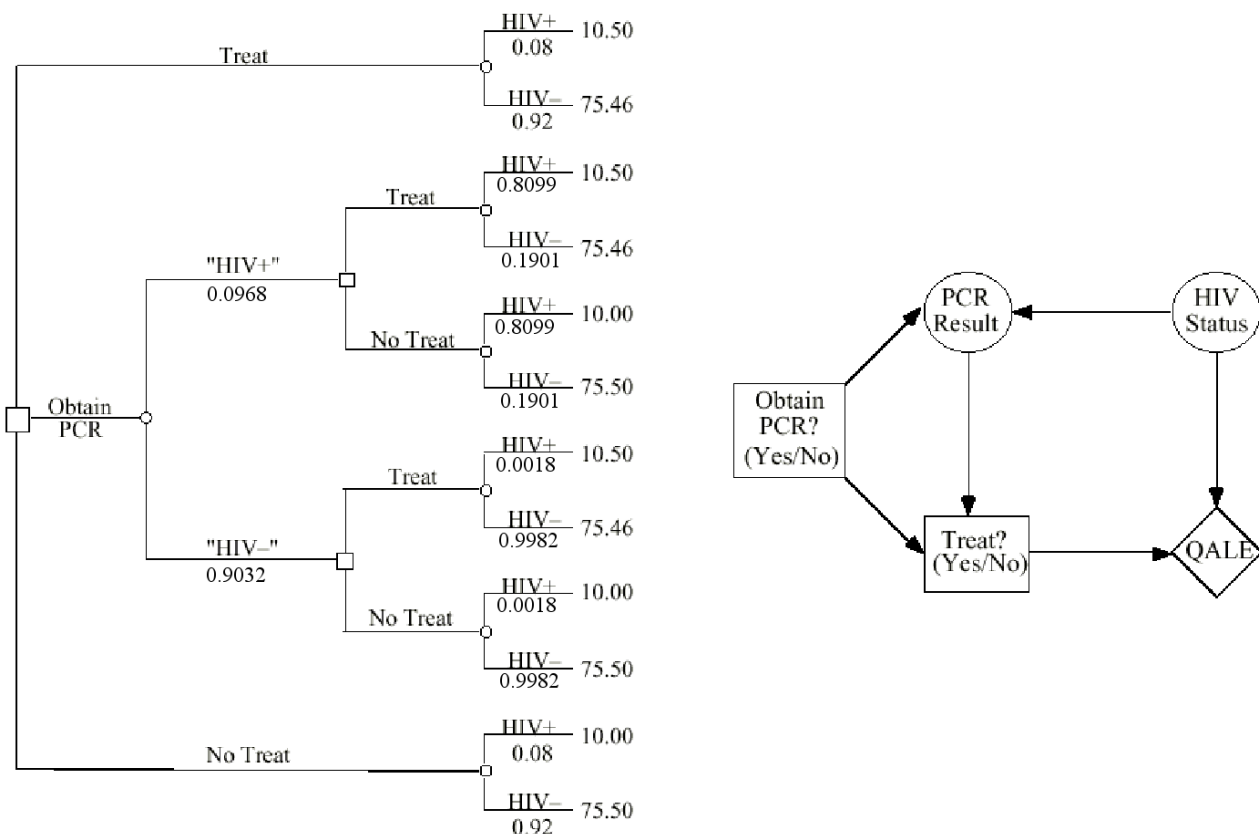


Representation of Decision Problems with Influence Diagrams

Influence diagrams are a powerful graphical representation for decision models, complementary to decision trees. Influence diagrams and decision trees are different graphical representations for the same underlying mathematical model and operations. So, if the two representations perform similar functions, why use an influence diagram? Influence diagrams represent the probabilistic structure of complex problems compactly, facilitate communication between analysts and decision makers, and form the basis for efficient and easy-to-use computer-based tools. However, influence diagrams also have limitations relative to decision trees.

1. Influence Diagrams: Definitions and Notation

Influence diagrams represent graphically many of the same components of a decision model as do decision trees. The following diagram shows a decision tree and the corresponding influence diagram. The decision problem modeled by the tree and by the influence diagram is the decision to test an infant born to a mother who is infected with the HIV (human immunodeficiency virus) with the PCR (polymerase chain reaction, a technique that is useful for HIV diagnosis in infants).



Both the decision tree and the influence diagram show two decisions. The first decision is whether to perform PCR, and the second is whether to treat once the test result is known. Test results are shown in quotation marks ("HIV+") and the true disease state is shown without quotation marks (HIV+). The payoff is measured in QALE - quality-adjusted life expectancy, a common method of expressing outcomes in medical decision making.

Two new graphical elements are apparent in the influence diagram: arcs between nodes, and the value node (shown as a diamond).

The *arcs* represent relationships between the *nodes*. A *decision node* (drawn as a rectangle) provides the decision alternatives under consideration. A *chance node* (drawn as a circle or oval) represents a variable whose value is a probabilistic function. An arc between two chance nodes indicates that a probabilistic relationship between the two events *might* exist. A probabilistic relationship exists when the occurrence of one of the events affects the probability of the occurrence of the other event. We know, for example, that the PCR test result (the chance node PCR Result in the previous diagram) depends on whether the infant is infected (HIV Status in the previous diagram). The arc between HIV Status and PCR Result indicates this probabilistic dependence. An arc represents a *weak* assertion about a probabilistic relationship, because an arc is allowable between two chance nodes when, in fact, no probabilistic relationship exists.

The arc points from the conditioning event to the conditioned event. Thus, in the previous diagram, the arc from HIV Status to PCR Result indicates that the test result is conditioned on the infection status; the diagram requires an assessment of the probability of a positive or negative test conditioned on whether the infant is infected. The direction of the arc, therefore, determines which probabilities will be assessed as conditional or as unconditional. The direction of arcs may be changed with Bayes' theorem during evaluation of an influence diagram. Therefore, the direction of the arc *does not* imply causality.

The absence of an arc is a *strong* assertion of independence or of *conditional independence*. Two events are conditionally independent, given a third event, if, after we have observed the third event, observing one of the two events gives us no additional information about the likelihood of the other event, i.e., A and B are conditionally independent, given C, when: $\Pr(A | C) = \Pr(A | B \& C)$.

In the previous diagram, there are arcs from HIV Status and Treat? to QALE, but no arc from PCR Result to QALE. This absence of an arc indicates that QALE is conditionally independent of PCR Result, given knowledge of HIV status and whether the patient was treated (Treat?). QALE is conditionally independent of PCR Result because, if we know whether the patient truly has HIV and whether she was treated, information about the test result, per se, does not affect quality-adjusted life expectancy. Arcs between chance nodes can be omitted only when events are assumed to be conditionally independent.

In contrast to a decision tree, in which the sequence of events is evident from the tree structure, an influence diagram relies on specific types of arcs to represent the sequence of events. An arc that points into a decision node from a chance node indicates that the chance event has been observed (or is known) at the time the decision is made. Thus, in the previous diagram, the arc from PCR Result to Treat? indicates that the decision maker knows the PCR result prior to making the decision to treat. Such arcs are called *informational arcs*. Conversely, the absence of an arc from a chance node (HIV Status) to a decision (Treat?) indicates that the decision maker has not observed the outcomes of the chance event when she makes the decision. Thus, the previous diagram asserts that the decision maker does not know HIV status when she decides whether to treat.

An arc that points from a decision node A (Obtain PCR?) into decision node B (Treat?) indicates that decision A is made prior to decision B. These arcs are often called *"no-forgetting arcs,"* indicating that the decision maker does not forget decisions that were made previously, or forget the information available at the time of the earlier decisions. (A decision tree requires these assumptions as well.) The

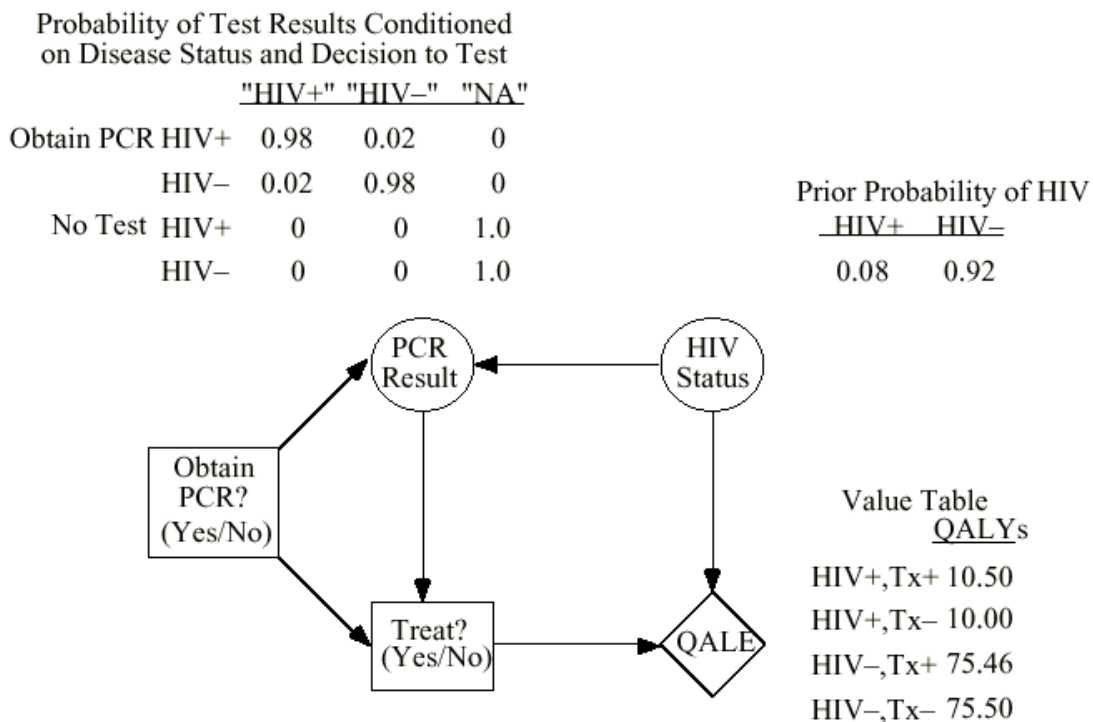
analyst must specify completely the order of decisions in the influence diagram; the no-forgetting arcs enable the analyst to indicate this ordering. In the above diagram, for example, the no-forgetting arc from Obtain PCR? to Treat? indicates that the decision maker decides whether to obtain PCR prior to deciding whether to treat and that the decision maker remembers whether she obtained a PCR test at the time she decides whether to treat.

The information associated with a node is determined by the node type and by the node's parents. The *parents* (or *direct predecessors*) of a node are those nodes that send arcs to the node; the *children* (or *direct successors*) of a node are those nodes that receive arcs from the node. A chance node provides the probability of the outcome of events conditioned on the node's parents. For example, in the above diagram, PCR Result provides the probability of a positive and of a negative test conditioned on HIV Status and on whether a test was ordered.

The node labeled QALE in the previous diagram is a *value node*. The value node, drawn as a diamond or sometimes as a hexagon, contains a table that represents the payoff (sometimes expressed as a utility) of all possible outcomes. The value node, in essence, contains the information that is shown at the ends of the branches of a decision tree.

The parents of the value node (i.e., nodes that have arcs that point to the value node) indicate the events and decisions that affect the payoff directly. The parents of the value node depend on the utility model used for the analysis. In the previous diagram, we see that only HIV Infection Status and Treat? affect value (quality-adjusted length of life) directly.

The following diagram shows the same influence diagram, but with tables adjacent to each of the nodes. These tables contain the information that would be associated with the branches in a decision tree. Two types of table are shown: probability tables and a value table. (Tx+ indicates the decision alternative "treatment offered" and Tx- "Treatment withheld").



The probability tables contain the probabilistic information that relates conditioned events and conditioning events. For example, the probability table associated with the PCR Result node in the previous diagram denotes the probabilistic relationship of the outcomes of the PCR Result chance node to the conditioning events, which are the decision Obtain PCR? and the HIV Status. The convention for these tables is that we place the conditioning events (Obtain PCR, HIV+) in a column at the left of the table, and place the possible values for the variable ("HIV+", "HIV-", "NA") across the top of the table. The probabilities in each row of the table sum to 1.0, because the row contains all possible outcomes conditioned on the events in the first column.

The probability table associated with PCR Result indicates that, if the decision alternative Obtain PCR? is chosen and if the patient is infected, the probability of a positive test result ("HIV+") is 0.98, and the probability of a negative test result ("HIV-") is 0.02. The table indicates that, if the decision alternative No PCR is chosen, then the test result is "not available" with probability 1.0 (right column), regardless of the infection status of the patient.

The need for probabilities for the test result when the No Test alternative is chosen arises because the influence diagram assumes symmetry of the alternatives and uncertain events.

When a decision problem has different events depending on the actions taken (for example, the events following the Obtain PCR and No PCR alternatives are different), we say that the problem has *structural asymmetry*. Thus, in the influence diagram, asymmetry in the structure of a decision is denoted by outcomes that have a probability of 1 or 0.

The value table is best understood in relation to the value node. The arcs into the value node (from the parents) indicate the events that affect the value of the outcomes. The value table is associated with the value node; it reflects how the decision maker values the possible outcomes that may result from the decision. In the first diagram, for example, the outcomes in the decision tree and the influence diagram are valued in terms of QALE (quality-adjusted life expectancy).

The arcs into the value node indicate that the QALE depends on the whether the screened person is infected (HIV Status), and on whether treatment is initiated. The value table indicates that treatment increases QALE by 6 months for infants infected with HIV. We assume that adverse effects from treatment result in a decrease of QALE of approximately 0.04 years for uninfected infants. Thus, in the influence diagram, the values of the outcomes are placed in the value table; in a decision tree, they would be placed at the end of the branches.

Examination of the two diagrams reveals several differences between the information displayed in the decision tree and that displayed in the influence diagram. First, the decision alternatives, the outcomes of the chance events, and the probabilities associated with these outcomes are not shown graphically in the influence diagram. In the influence diagram, this information is contained in tables associated with the corresponding nodes (as in the previous diagram).

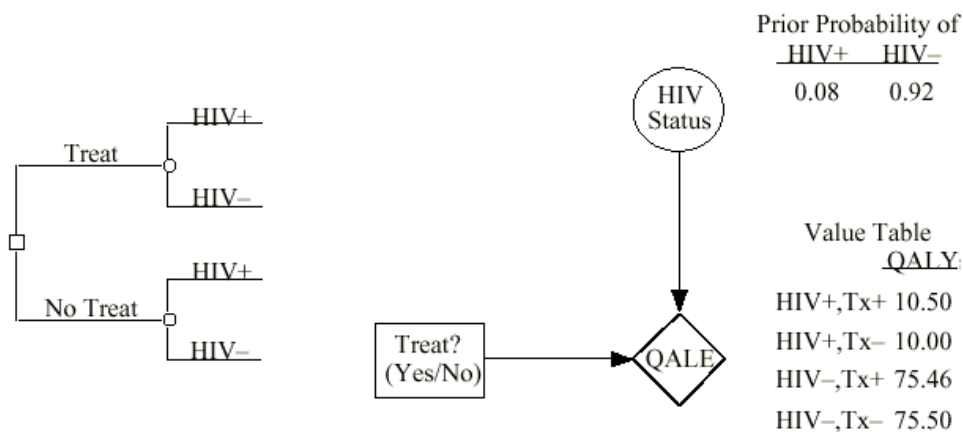
In software developed to analyze influence diagrams, data entry is usually made directly into such tables. The emphasis on the probabilistic structure of the problem in the graphical representation of influence diagrams results in the enhanced capability for displaying the probabilistic relationships in large, complex problems; however, the probabilities and utility of outcomes are not apparent in the graphical representation. Note that the decision tree shows structural asymmetry (for example, the Obtain PCR branch and the Treat branch of the decision tree in the first diagram are different) graphically, and that the influence diagrams highlight the probabilistic relationships in the model. For

example, the influence diagram indicates that QALE is independent of the test result, given knowledge of the infection status, and of whether treatment was initiated.

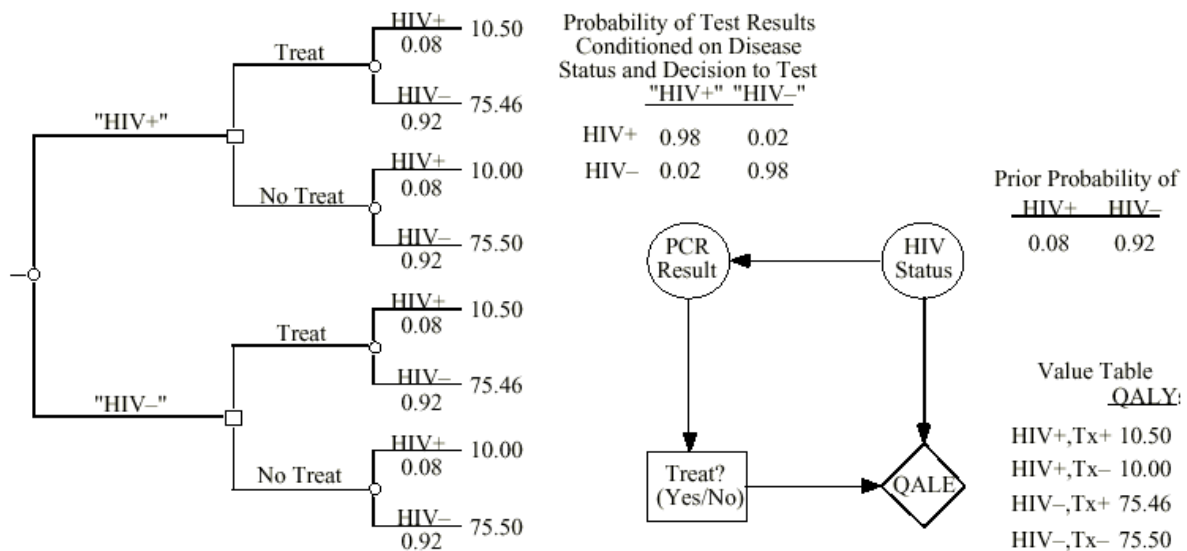
The absence of an arc from Obtain PCR? to QALE indicates that, conditioned on the disease state and treatment decision, the test is assumed to have no effect on outcome (a harmless test). Assumptions about probabilistic independence are explicit in the influence diagram, even in extremely complex models. Such assumptions are more difficult to identify with a decision tree: The analyst must compare the probabilities on all branches in the tree for the relevant events. A decision tree shows structural asymmetry in branches that model different events and, thus, differ in their structure. In the influence diagram, however, structural asymmetry is hidden; the analyst indicates asymmetry by assigning a probability of 0 to certain events. For problems with considerable structural asymmetry, the decision tree may be a more natural representation.

2. Influence Diagrams and Trees: More Examples

The following tree and influence diagram show a treatment decision for HIV infection in the infant when test information is not available. From the influence diagram, we see that the value depends on whether treatment is offered and on whether HIV infection is present.

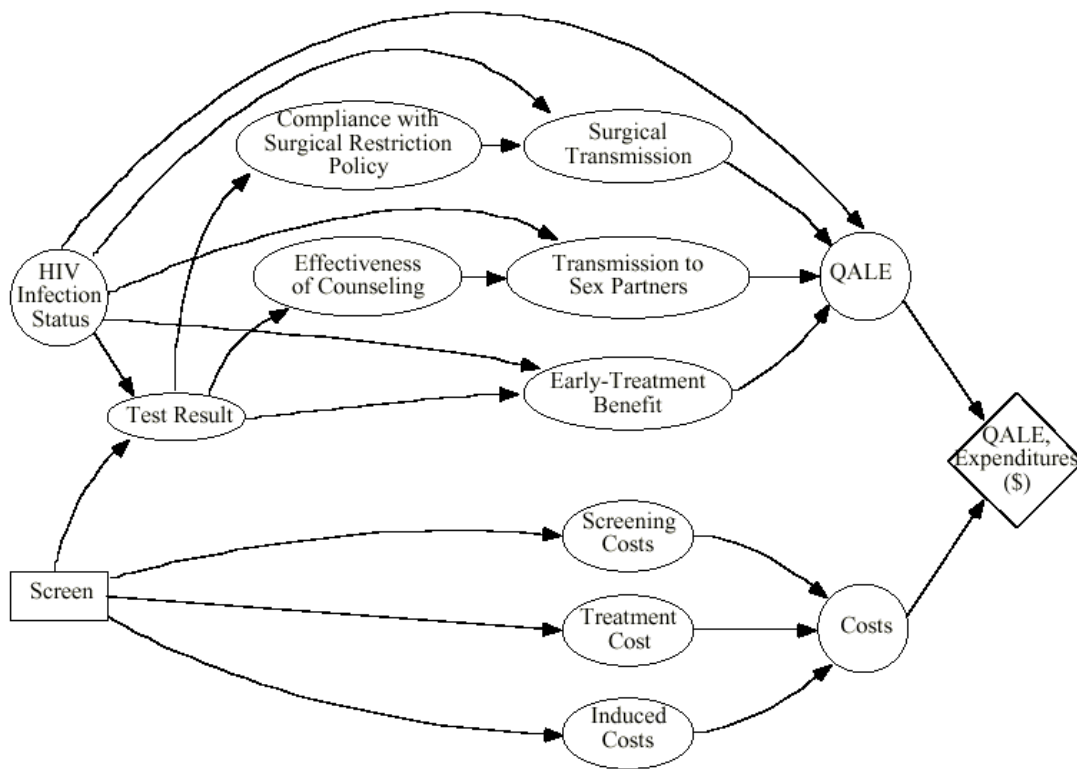


The next tree and influence diagram represent a treatment decision in which test information is available.



The decision tree begins with a chance node, because we assume that the decision maker has already ordered a test, that she has not yet received the result (therefore, the result is uncertain), and that she will have observed the test result when she makes the decision to treat. From the influence diagram, we note again that the value depends on whether treatment is offered and on whether disease is present. The PCR Result node is probabilistically related to the HIV Status node, as indicated by the arc from HIV Status to PCR Result. The informational arc from PCR Result to the Treat? decision node indicates that the decision maker will have observed the test result at the time that she makes the decision. The absence of an arc between HIV Status and Treat? indicates that the decision maker does not know the disease status at the time that she makes the decision, and that her decision does not affect Disease Status.

For the simple problems described above, the decision-tree representation is appealing. For large and complex problems, however, representation in a tree becomes more challenging. The following diagram shows an influence diagram for a model of a cost-effectiveness analysis (measuring cost per quality-adjusted life year) of a program to screen surgeons for HIV to prevent transmission to patients.



Working back from the value node, we see that the value node is affected directly by QALE and by Costs. Examining the parents of QALE, there are three benefits of a screening program: the benefit from reduced transmission to patients during surgery (from a policy that restricts surgeons identified as having HIV from performing procedures), the benefit from reduced transmission of HIV to the surgeons' sexual partners, and the benefit to the surgeon screened and identified as having HIV (from early medical intervention). The degree to which these benefits are realized depends in turn on compliance with the policy to restrict surgical practice, and on the effectiveness of counseling in inducing reductions in sexual risk behaviors. The influence diagram shows that each of these benefits depends on the test result, which depends on infection status, and the decision to screen. Costs are modeled in a similar manner. Conditional independence is shown by the absence of arcs.

3 Evaluation of Influence Diagrams

Because decision trees and influence diagrams are differential graphical representations of similar underlying mathematical operations, it is possible to convert one graphical representation into the other. Any symmetric decision tree can be drawn as an influence diagram without further manipulation; however, all influence diagrams cannot be drawn as decision trees directly.

Decision trees usually are drawn from left to right, with events in the *order of observation relative to the decisions*. Thus, an event that is observed by the decision maker before she makes a decision is drawn to the left of the decision; an event that is observed after the decision is drawn to the right of the decision. To draw the tree in order of observation, the analyst may need to use Bayes' theorem to calculate the probabilities required in the tree (the probabilities of a positive and of a negative test result, and the post-test probability of disease given each test result).

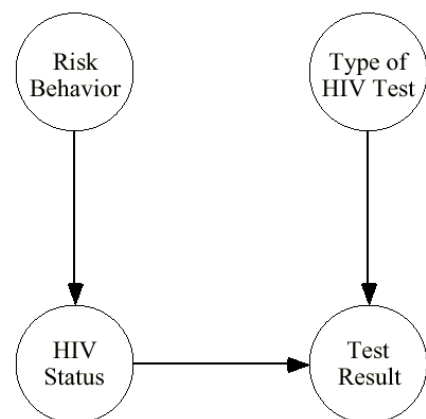
Because an influence diagram does not require events to be drawn in the order of observation, the analyst may need to manipulate the diagram before converting it to a decision tree. This manipulation is mathematically identical to the use of Bayes' theorem to calculate probabilities required in the decision tree. Applying Bayes' theorem in an influence diagram involves changing the direction of the arcs, an operation called *arc reversal*, and an essential operation for the evaluation of an influence diagram.

3.1 Arc Reversal.

An influence diagram can be drawn to reflect either of two different orders of probabilistic conditioning. We can draw the diagram with variables ordered according to when they will be observed (as in a decision tree). One of the advantages of the influence diagram, however, is that it can be drawn with the variables ordered such that the probabilities can be assessed most easily. This order of conditioning is called the *assessment ordering*. If, for example, we have evidence about the prevalence (prior probability) of a disease and about the sensitivity (the likelihood that the presence of the disease will be correctly detected) and specificity (the likelihood that the absence of the disease will be concluded correctly) of a diagnostic test, then the assessment ordering is disease status (e.g., HIV Status), followed by the test result (e.g., PCR Result). The ability to use assessment ordering freely can ease the analytic task substantially. It allows the analyst to build the diagram with events in the order that most facilitates probability assessment (for example, from cause [disease status] to effect [test result]), an approach to structuring the problem that is more natural than the reverse ordering.

We can illustrate the concepts underlying arc reversal with an example. Suppose that we want to determine the probability that a person who has a history of injection drug use and needle sharing has HIV. Let's assume that two tests are available: Test1, which has a both a sensitivity and a specificity of 98.1, and Test2, which have a sensitivity of 99.5% and a specificity of 99.99%. The influence diagram alongside illustrates what we know.

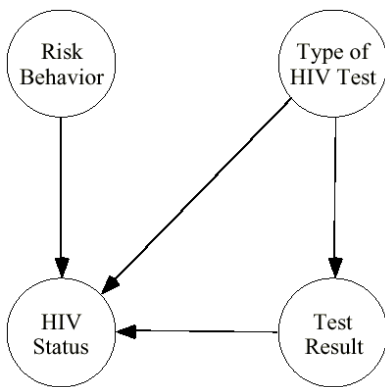
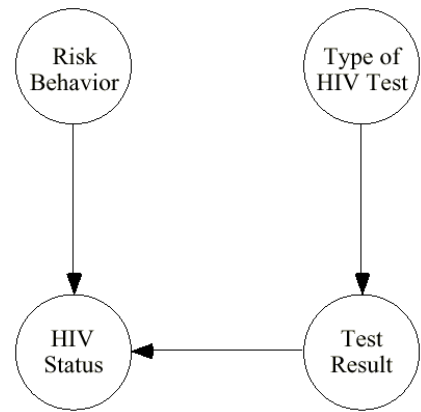
We believe that HIV Status depends on Risk Behavior — in this case, needle sharing. More specifically, the prior probability of HIV depends on risk behavior. Test Result depends on Type of HIV Test (because one test is more accurate than the other, and



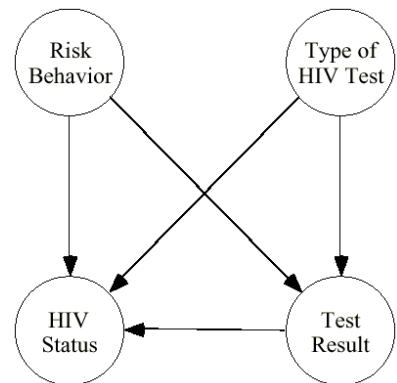
thus the sensitivity and specificity will differ), and on HIV Status. Because Test Result is conditioned on HIV Status (rather than HIV Status being conditioned on Test Result), the diagram is drawn in assessment ordering, and we can enter the sensitivity and specificity of each test directly in the probability table associated with the node Test Result.

The diagram asserts that Test Result is independent of Risk Behavior, given that we know HIV Status. This independence is indicated by the absence of an arc from Risk Behavior to Test Result, and indicates that if we knew the true HIV status of an individual, information about their risk behaviors would not affect the probability of a positive or negative test result. We also note from the diagram that Risk Behavior and HIV Status are independent of Type of HIV Test, as shown by the absence of arcs between these nodes. We could complete the probability tables in this influence diagram with information that we already know — the prevalence of HIV conditioned on risk behavior (for HIV Status) and the sensitivity and specificity of the HIV tests (for Test Result). The information we seek from the influence diagram, however, is the post-test probability of HIV infection, given a particular test result. That is, what is HIV Status, given Test Result?

We can answer this question by reversing the arc between HIV Status and Test Result, as shown alongside; this diagram indicates that HIV Status is conditioned on Test Result. The probability table associated with HIV Status would therefore contain the probability of HIV conditioned on a positive or negative test. But is this diagram otherwise correct? It asserts that HIV Status depends on only Test Result (that is, whether the test result is positive or negative), and is independent of Type of HIV Test. We know, however, that the posterior probability of HIV depends on which test we choose because the tests have different sensitivities and specificities. Thus, simply reversing the arc between HIV Status and Test Result creates a diagram that incorrectly represents our knowledge of the problem.



To correct this problem, we first add an arc from Type of HIV Test to HIV Status, as shown here. This new arc reflects the dependence of the posterior probability of HIV infection on the type of test that we choose. Is this revised diagram correct? It asserts that Test Result depends on only Type of HIV Test, and is independent of Risk Behavior. However, we know that Test Result depends on Risk Behavior; a positive test result is more likely among those patients who engage in risky behavior.



To reflect this dependence, we must add an arc from Risk Behavior to Test Result, as shown here. Note that now both HIV Status and Test Result are conditioned on the same events. We now have a diagram that represents our knowledge faithfully, and because the HIV Status is conditioned on Test Result, we should be able to determine the posttest probability of HIV infection.

The reasoning about this example illustrates a general principle. To reverse an arc between node A and node B, we must ensure that node A and node B have the same parents — A and B must be conditioned on the same events. Often, as we did in the above example, we must add arcs to the diagram to meet this criterion. The reason is that mathematically, we perform arc reversal using Bayes' theorem. Bayes' theorem requires that events A and B be conditioned on the same events. In summary, we can draw the influence diagram such that we facilitate data entry, and we can use arc reversal to perform inference.

3.2 The Evaluation Algorithm

The mathematical operations used to evaluate influence diagrams and decision trees are the operations of probability and expected-utility theory. The key concept is that decision trees and influence diagrams provide different graphical representations of the same underlying probability distributions and expected utility operations.

The procedure for evaluating the decision tree can be performed easily by hand. Starting at the right side of the tree, the analyst takes the expectation at each chance node, substitutes the expectation for the chance node, and repeats the process until she reaches a decision node. Thus, the analyst removes the nodes by taking the weighted average of the branches of the node. This process is called ***folding back*** or ***chance-node removal by averaging***. At the decision node, the analyst chooses the alternative with the highest expected utility, and then removes the decision node. This process is called ***pruning*** or ***decision-node removal by policy determination***. The tree is successively whittled away and simplified as it is evaluated. This process continues until the analyst can determine the expected utility of the decision alternatives at the leftmost decision.

Evaluation of an influence diagram is similar — the analyst simplifies the diagram successively until she can evaluate expected utility of the decision alternatives. The process for simplifying the diagram also makes use of chance-node removal by averaging and decision-node removal by policy determination. Because the analyst may have drawn the influence diagram in assessment ordering, she may need to reverse arcs (and therefore to add arcs) before she can remove a node by averaging.

There are several algorithms for evaluating influence diagrams. The simplest is the method of arc reversal and node removal, developed by Shachter. Although small influence diagrams can be evaluated by hand, in practice, software is required, much as it is for complex decision trees or for those trees that require substantial Bayesian updating.

We must remove nodes from an influence diagram in a specific order, just as we must evaluate the nodes in a decision tree in a specific order. The general principles are similar to those used to evaluate a decision tree. We first remove chance nodes for events whose outcomes are revealed, if ever, subsequent to a decision — these are events that the decision maker has not observed at the time of the decision. We remove decision nodes in reverse of the order in which we will actually make the decisions.

That is, we remove the decision node for the final decision first, in a manner analogous to the evaluation of a decision tree that has sequential decisions. After removing the decision node for the final decision, we remove chance nodes whose events are revealed subsequent to the next-to-final decision that we must make, and so on.

We can evaluate influence diagrams with the following algorithm that formalizes these concepts:

1. Eliminate all nodes (except the value node) that do not point to another node (*barren nodes*). They may occur when several types of evidence are observed in large models.
2. As long as there are one or more nodes that point into the value node, do the following:
 - a. If there is a decision node that points into the value node, and if all other nodes that point into the value node also point into that decision node, remove the decision node by policy determination. Remove any nodes (other than the value node) that no longer point to any other node. Go back to step 2.
 - b. If there is a chance node that points into only the value node, remove it by averaging. Go back to step 2.
 - c. Find a chance node that points into the value node and not into any decision, such that, if you reversed one or more arcs pointing from that chance node (without creating a cycle), the chance node would point into only the value node. Reverse those arcs (adding arcs as needed). Go back to step 2.

After a node of any type is removed, draw the arcs from its parents to the value node.

Step 2a recognizes that if all chance nodes point into a decision node (that is, there are informational arcs from all chance events into a decision), then we have observed the outcomes of all chance events, and so have no uncertainty about the relative ranking of the decision alternatives. We can, therefore, choose the best alternative, and remove the decision node by policy determination.

Step 2b recognizes that if a chance node points into the value node and into no other decision nodes, then the outcomes of the uncertain event are revealed subsequent to any decisions, so we can remove the chance node by averaging. Removing nodes by averaging is analogous to folding back the tree from the endpoints of the branches in the tree.

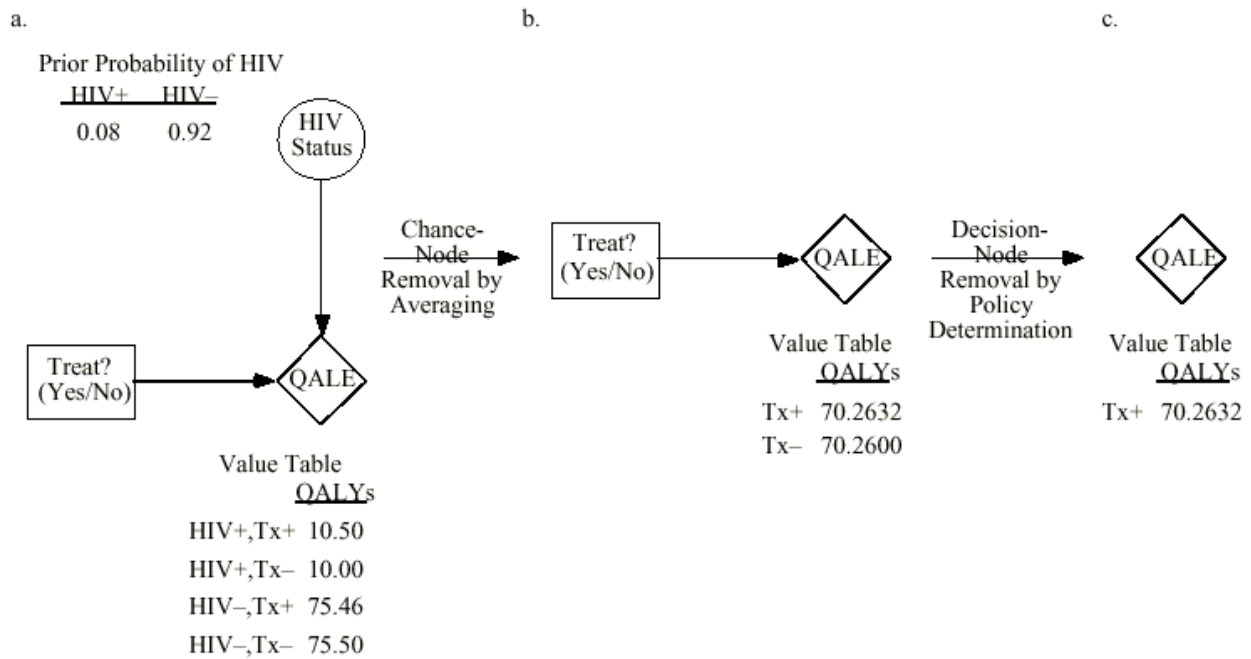
Step 2c identifies chance nodes whose events are revealed subsequent to all decisions (these nodes will not have informational arcs into any decision node, which indicates that their outcomes are not observed prior to the decisions), but that must undergo arc reversal prior to removal. These are variables whose order must be reversed if we are perform the evaluation. The conditions addressed by step 2c occur only when we have entered probabilistic information into the influence diagram such that the application of Bayes' theorem is required during evaluation. The operation of arc reversal is analogous to the use of Bayes' theorem, which is often used in the process of building a decision tree (prior to evaluation of the tree).

4 Influence Diagram Evaluation – Worked Examples

4.1 The Treat or No Treat Decision.

The decision maker must decide whether to treat or not to treat, with no test information. There are no barren nodes (step 1). The conditions of step 2a are not satisfied — HIV Status does not point to the decision Treat?, which indicates that the decision maker has not observed HIV Status at the time that she makes the decision — the decision node cannot yet be removed.

The conditions of step 2b are satisfied, however, so we remove HIV Status by averaging:



To determine what calculation to perform in removing a node, we redraw the diagram with the arcs from the parents of the removed node pointing to its children, then examine the diagram to determine the appropriate conditioning for the new probability or value tables. If the node we wish to remove has no parents, as is the case if part a of the diagram, we simply remove the node, and then examine the diagram to determine the appropriate conditioning. We see from part b of the diagram that, after we remove HIV Status, QALE will be conditioned on only Treat?. Therefore, we must calculate the expected utility for each of the decision alternatives for the Treat? decision using the information in the probability table associated with HIV Status and in the value table associated with QALE. Note that performing this operation is identical to folding back the first decision tree on page 5 one level and removing the chance node from the tree. We calculate the expected utility for the treatment option (Tx+) first:

$$\begin{aligned}
 EU[Tx+] &= p(\text{HIV+})U[\text{HIV+}, Tx+] + p(\text{HIV-})U[\text{HIV-}, Tx+] \\
 &= (0.08)(10.50 \text{ QALYs}) + (0.92)(75.46 \text{ QALYs}) \\
 &= 70.2632 \text{ QALYs.}
 \end{aligned}$$

The expected utility for the no treatment option (Tx-) is

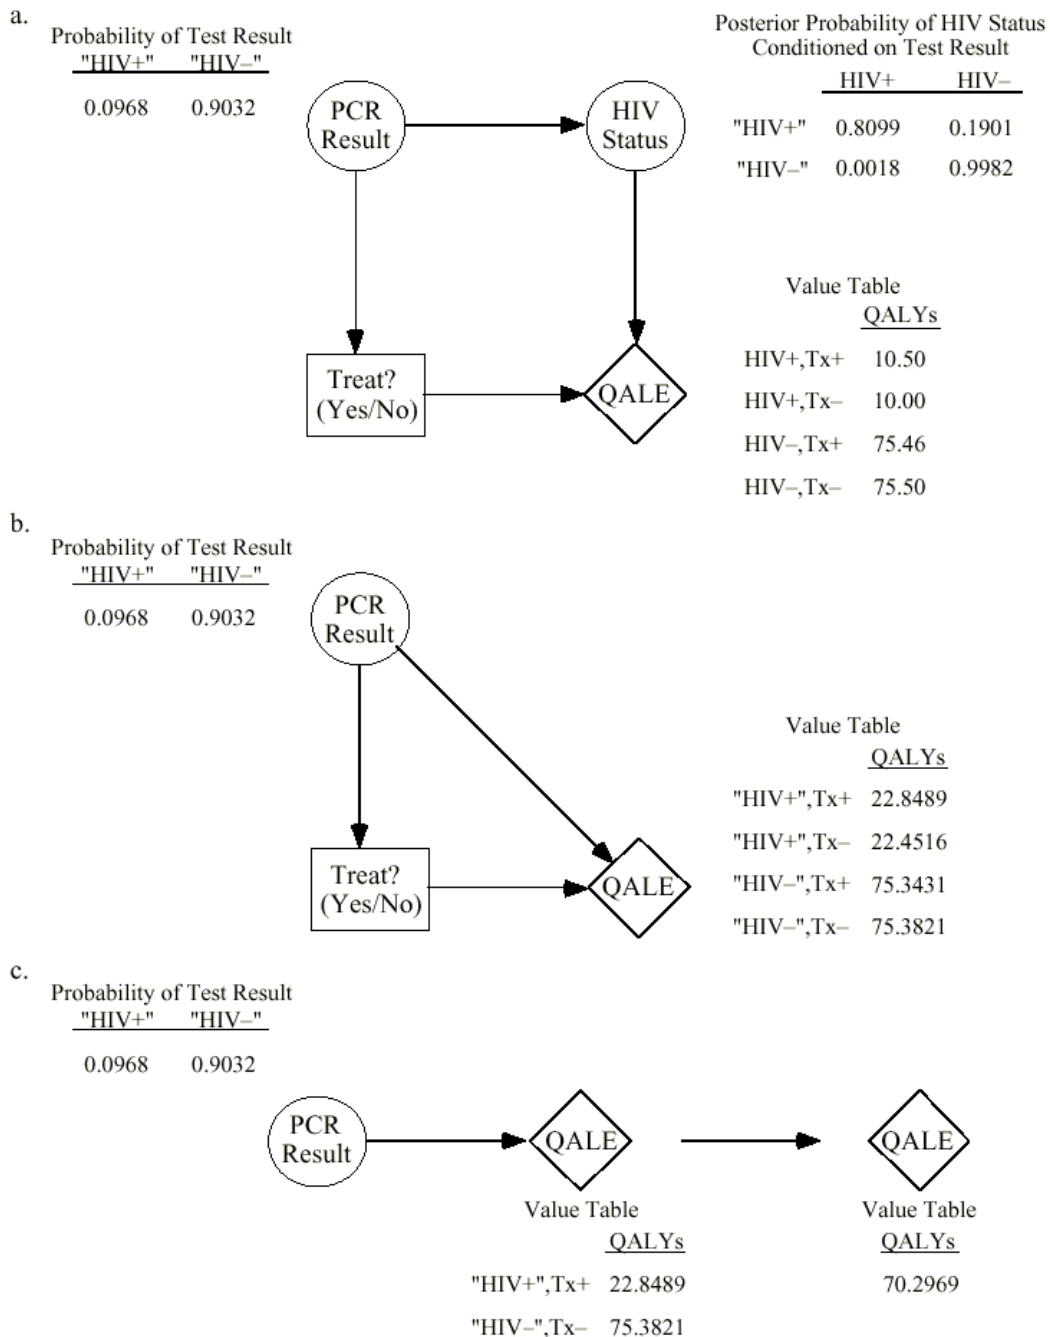
$$\begin{aligned}
 EU[Tx-] &= p(\text{HIV+})U[\text{HIV+}, Tx-] + p(\text{HIV-})U[\text{HIV-}, Tx-] \\
 &= (0.08)(10.00 \text{ QALYs}) + (0.92)(75.50 \text{ QALYs}) \\
 &= 70.2600 \text{ QALYs.}
 \end{aligned}$$

We now have accounted for our uncertainty about HIV status, and have calculated the expected utility associated with each of our decision alternatives; we can therefore remove the node HIV Status (part b of the diagram).

Continuing to follow the algorithm, we return to step 2. We have removed by averaging all the chance nodes whose events are revealed subsequent to the decision, so we can choose the best policy. That is, the diagram satisfies the conditions in step 2a, and we can remove the decision node by policy determination (part c of the diagram). The Treat alternative has a higher quality-adjusted life expectancy, but the advantage is small. We now have evaluated this influence diagram fully, and have determined that the Treat alternative has the higher expected utility (70.2632 QALYs).

4.2 The Treat or No Treat Decision with a Diagnostic Test.

We now extend the previous example to include diagnostic-test information (for this example, we start with the second influence diagram on page 5, and show how to evaluate this diagram:



We assume, for this example, that we have already ordered a diagnostic test, and that we shall know the test result at the time we make the Treat? decision. We see that there are no barren nodes (Step 1). We cannot yet make the Treat? decision, because the outcomes of the chance event (HIV Status) are revealed subsequent to the decision (we know that because there is no informational arc from HIV Status to Treat?). Thus, the conditions of step 2a are not satisfied. We cannot remove HIV Status by averaging because it points to PCR result, thus failing the conditions in step 2b. We note that, however, if we reverse the arc between HIV Status and PCR Result, the conditions in step 2c are satisfied — we can remove HIV Status by averaging. As shown in part a of the diagram, we can reverse the arc between PCR Result and HIV Status without adding arcs, because neither node has other parents.

Note how the information in the probability tables has changed when the diagram is changed from assessment order (the second influence diagram on page 5) to observation order (part a of the previous diagram). The probability table for PCR Result originally contained the sensitivity and specificity of PCR, and the table for HIV Status contained the prior probability of HIV infection (the second influence diagram on page 5). In part a of the previous diagram, we used Bayes' theorem to reverse the arc and to calculate the probability of a positive or negative test (the denominator of Bayes' theorem provides the overall probability of a positive or negative test), as shown in the probability table for PCR Result, and the posterior probability of HIV conditioned on positive and negative test result, as shown in the probability table for HIV Status.

After reversing the arc, we return to step 2 of the algorithm. We note that we cannot make the Treat? decision because the outcomes of the chance events represented by HIV Status are revealed subsequent to the Treat? decision; thus, the conditions in step 2a are not satisfied. The conditions of step 2b are satisfied: HIV Status now points to only the value node. Therefore, we can remove HIV Status by averaging; that is, we calculate the expected utility for each of the decision alternatives conditioned on the test result. Now, however, HIV Status is conditioned on PCR Result, so when we remove HIV Status, we must draw an arc from PCR Result to QALE (part b of the previous diagram). Thus, after we remove HIV Status, QALE will still be conditioned on PCR Result and on Treat? We must therefore calculate the entry in the value table for each possible combination of QALE's parents: "HIV+", Tx+; "HIV+", Tx-; "HIV-", Tx+; and "HIV-", Tx-.

We use the probabilities from the probability table associated with PCR Result to calculate, for example, the entry in the value table for "HIV+", Tx+ as

$$\begin{aligned} EU["HIV+", Tx+] &= p(HIV+|"HIV+")U[HIV+, Tx+] + p(HIV-|"HIV+")U[HIV-, Tx+] \\ &= (0.8099)(10.50 \text{ QALYs}) + (0.1901)(75.46 \text{ QALYs}) \\ &= 22.8489 \text{ QALYs} \end{aligned}$$

Similarly, we can calculate the entry in the value table for "HIV+", Tx- as

$$\begin{aligned} EU["HIV+", Tx-] &= p(HIV+|"HIV+")U[HIV+, Tx-] + p(HIV-|"HIV+")U[HIV-, Tx-] \\ &= (0.8099)(10.00 \text{ QALYs}) + (0.1901)(75.50 \text{ QALYs}) \\ &= 22.4516 \text{ QALYs} \end{aligned}$$

The entries in the value table in part b of the previous diagram show the QALE for each of the strategies.

We return to the algorithm and note that the conditions for step 2a are now satisfied: The decision node Treat? points into the value node, and the only other node that points into the value node (PCR status) also points into Treat?. Thus, we can remove Treat? by policy determination. Unlike in our previous example, however, at the time that we make the decision, we will have observed a PCR test result. Therefore, we shall continue the example by allowing for either a positive or a negative result to occur, and evaluating the best treatment option conditioned on the test result. We see by inspection of the value table in part b of the previous diagram that, as expected, if the PCR result is positive, we obtain a higher quality-adjusted life expectancy by choosing to treat (22.8489 QALYs) than by choosing not to treat (22.4516 QALYs), and if the PCR result is negative, we obtain a higher quality-adjusted life expectancy by choosing not to treat.

The resulting diagram is shown in part c of the previous diagram. Because we continue to show PCR Result as a chance node, the diagram reflects that we have not yet observed the result of the PCR test, which could be either positive or negative. If the test is positive, the expected utility of our decision (treat) is 22.8489 QALYs; if the test is negative, the expected utility of our decision (no treatment) is 75.3821 QALYs. We can calculate the quality-adjusted life expectancy for all patients who undergo testing by following the algorithm and removing the chance node by averaging, as shown in part c of the previous diagram. We multiply the probability of a positive PCR test result by the expected utility given a positive result, and add this value to the product of the probability of a negative test result and the expected utility given a negative test result, or

$$\begin{aligned} & p(\text{"HIV+"})EU[\text{"HIV+"}, T_{x+}] + p(\text{"HIV-"})EU[\text{"HIV-"}, T_{x-}] \\ &= (0.0968)(22.8489 \text{ QALYs}) + (0.9032)(75.3821 \text{ QALYs}) \\ &= 70.2969 \text{ QALYs.} \end{aligned}$$

This result indicates that the quality-adjusted life expectancy for an entire cohort of patients (with a prior probability of HIV infection of 0.08) who undergo testing is 70.2969 QALYs.

In summary, to evaluate this diagram, we performed arc reversal using Bayes' theorem, removed the chance node by averaging, and then removed the decision node by policy determination. This series of steps left us with a diagram that showed the expected utility conditioned on the PCR test result. In practice, of course, only one result would occur, and we would observe it before making the Treat? decision. We finished the evaluation by removing this final chance node.

4.3 The Treat or Test Decision.

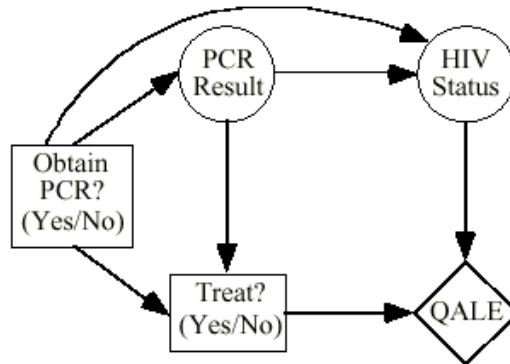
We now return to our opening example, the treat-or-test decision (page 3), and begin the evaluation as we did in the previous example. We cannot remove any nodes directly, so we reverse the arc between PCR Result and HIV Status to remove the node by averaging. To reverse this arc, we must add an arc from Obtain PCR? to HIV Status, so that each node has identical parents (part a of the next diagram). We calculate the probability of possible test results (positive, negative, or not available), and the posterior probability of HIV conditioned on each result, as in the previous example. If the test is not ordered, the posterior probability for HIV is the same as the prior probability, as shown in the probability table for HIV Status. We remove HIV Status (part a of the next diagram), and by averaging calculate the entries in the value table conditioned on PCR Result and Treat?. Referring to the algorithm, we note that we can now remove Treat? by policy determination (part c of the next diagram). To do so, we recalculate the value table conditioned on the alternatives for Obtain PCR?, and on PCR Result (part a of the next diagram). Referring to the algorithm again, we note that we can now remove PCR Result.

a. Probability of PCR Result

	"HIV+"	"HIV-"	"NA"
Obtain Test	0.0968	0.9032	0
No Test	0	0	1

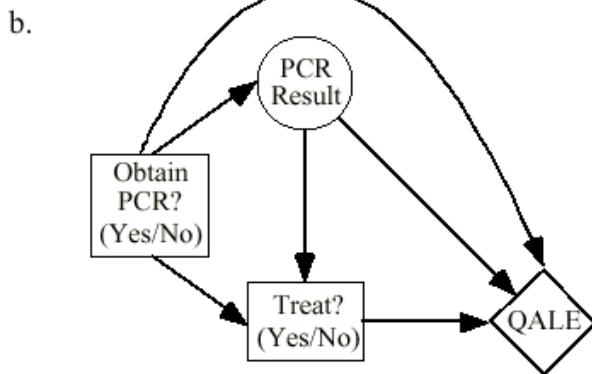
Posterior Probability of HIV Status Conditioned on Test Result

Result	HIV+	HIV-
Obtain Test "HIV+"	0.8099	0.1901
"HIV-"	0.0018	0.9982
No Test "NA"	0.08	0.92



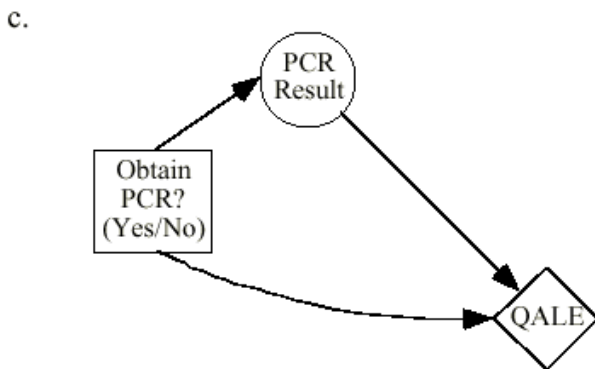
Value Table

	QALYs
HIV+,Tx+	10.50
HIV+,Tx-	10.00
HIV-,Tx+	75.46
HIV-,Tx-	75.50




Value Table

	QALYs
Obtain Test "HIV+",Tx+	22.8489
"HIV+",Tx-	22.4516
"HIV-",Tx+	75.3431
"HIV-",Tx-	75.3821
No Test "NA",Tx+	70.2632
"NA",Tx-	70.2600



Value Table

	QALYs
Obtain Test "HIV+"	22.8489
"HIV-"	75.3821
No Test "NA",Tx+	70.2632



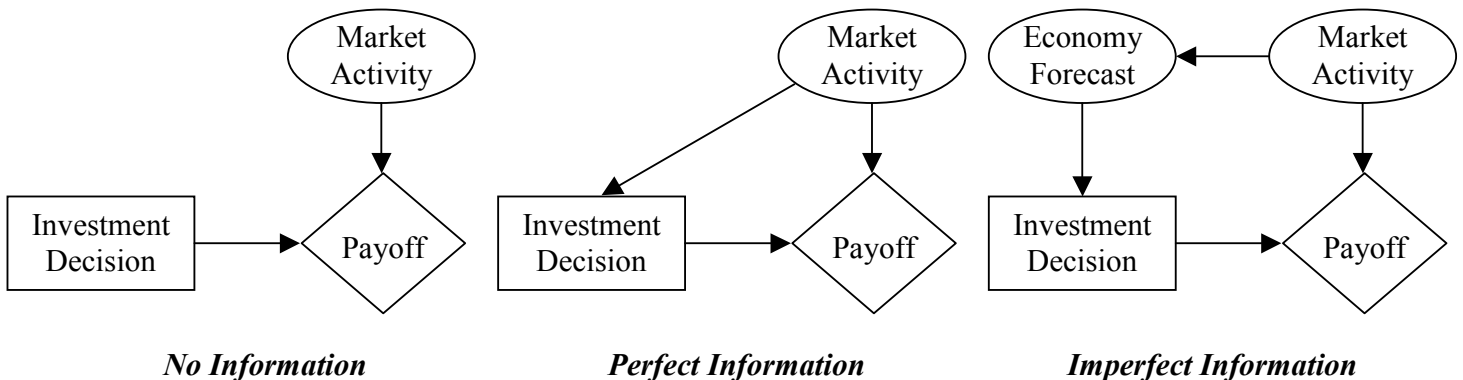
We calculate the expected utility for Obtain PCR alternative by averaging as

$$\begin{aligned}
 & p(\text{"HIV+"})EU[\text{"HIV+", Tx+}] + p(\text{"HIV-"})EU[\text{"HIV-", Tx-}] \\
 & = (0.0968)(22.8489 \text{ QALYs}) + (0.9032)(75.3821 \text{ QALYs}) \\
 & = 70.2969 \text{ QALYs.}
 \end{aligned}$$

Note that this calculation is the same as the calculation that we performed at the end of the previous example. The expected utility of the No Test alternative is 70.2632 (calculated in the example in Section 4.1, part c), as noted in the value table in part c of the previous diagram. We now can evaluate the decision Obtain PCR? and remove the decision node by policy determination (part d of the previous diagram). We see that the expected utility of the Test alternative is higher than that of the No Test (and Treat) alternative, but by only 0.0337 QALYs, which suggests that the decision is a close call. The decision is a close call despite the high sensitivity and specificity of PCR, primarily because in the absence of the opportunity to test, the optimal alternative is to treat all infants born of HIV-infected mothers, and testing only prevents unnecessary treatment of uninfected infants (which reduces quality-adjusted life expectancy of the uninfected infants by only 0.04 QALYs).

5 Using Influence Diagrams to find the Value of Information

Influence diagrams provide a really neat way of handling value of information problems, because information available for a decision can be represented through appropriate use of arcs and, if necessary, additional chance nodes representing imperfect (sample) information. The expected value of (either perfect or sample) information is simply the difference between the calculated expected value of the diagram with the informational arc and the calculated expected value of the diagram without information.



6 Strengths and Limitations of Influence Diagrams

Influence diagrams are particularly helpful when problems have a high degree of conditional independence, when compact representation of extremely large models is needed, when communication of the probabilistic relationships is important, or when the analysis requires extensive Bayesian updating.

Influence diagrams represent the relationships between variables. These relationships are important because they reflect the analyst's, or the decision maker's, knowledge about a problem. The construction of such a model often involves collaboration between an analyst and the decision maker.

This collaboration represents an exercise in knowledge acquisition — the analyst attempts to construct a model that reflects the decision maker's understanding of the problem domain. For example, in building the diagram on page 6, the analyst might first ask an expert for the major variables that affect health outcomes (Surgical Transmission, Transmission to Sexual Partners, and Early-Treatment Benefit). The analyst could then discuss the probabilistic relationship among the variables, and the factors that, in turn, affect each of the identified variables. By building the graphical elements of the influence diagram, the analyst can focus on the relationships among variables, before adding the detail needed for the associated probability and value tables. For problems in which the probabilistic relationships are complex, or have special importance, influence diagrams may be a useful aid in this knowledge acquisition task.

Another feature of influence diagrams is that they can be drawn with conditioning displayed in the manner that most facilitates assessment of the probabilities. This feature may also facilitate knowledge acquisition. For example, a clinician expert may be able to assess the prevalence of disease and sensitivity and specificity of a diagnostic test more easily than she could assess the post-test probability of disease. After the influence diagram is drawn to facilitate probability assessments, all updating and Bayesian inference are handled automatically by the evaluation algorithms. Although there are approaches for performing Bayesian updating within a decision tree, for problems with extensive Bayesian updating, such as sequential-testing decisions, influence diagrams ease the burden on the analyst by reducing the need for complex equations required for Bayesian updating in the tree. Influence diagrams also reduce the time required to find errors that may be introduced when these equations are specified.

Although influence diagrams offer advantages for certain analytic problems, they also have limitations relative to the decision-tree format. Highly asymmetric problems may be easier to understand when represented as decision trees. The timing of events may be easier to identify in a decision tree, although the same information is explicit in the influence diagram. In addition, it is not enough to simply draw the influence diagram. To represent fully the decision alternatives, strategies, alternative events, and value of outcomes, the analyst must complete the probability and value tables in the influence diagram, a process similar to placing the value of variables in the decision tree. As with a decision tree, an influence diagram for a complex problem may require a large number of probability assessments. For an asymmetric decision problem, the analyst must use probabilities of 0 or 1 to represent the asymmetry. The size and complexity of the probability and value tables also tend to increase rapidly as the problem modeled becomes more complex. The probability tables also become substantially larger if the chance events have multiple outcomes.

The evaluation algorithms for influence diagrams are designed for computer-based implementation. A moderate-sized tree can be solved by hand; only the simplest of influence diagrams could be solved readily without software. As arcs are reversed during evaluation of complex diagrams, extra arcs are added and the diagram may become confusing.

The choice of graphical representation should be governed by convenience, and will depend on the problem being analyzed, on the experience of the analyst, and on the background of the decision makers.