Decision tree analysis: Drawing some of the uncertainty out of decision making

William E. Marsh, PhD

Summary *— In veterinary practice, treatment outcomes and their economic consequences are often uncertain.* Decision analysis *(a formal, structured approach to making decisions when uncertainty exists) is applied to two situations commonly encountered in swine practice: whether or not to treat a sick animal, and how to determine the financial worth of using ultrasound to test for pregnancy in swine herds under various scenarios. Once a decision analysis model is developed, practitioners can use production record data from a client's herd, plus some knowledge of the characteristics of the test, as well as an estimate of the prevalence of the condition in the herd to customize the tree to illustrate the benefit of a certain decision.*

Most decisions involve some degree of uncertainty. We must constantly choose from several possible courses of action without knowing what the outcome of our decision will be. Agriculture is a particularly risky business. Farm managers are constantly at the mercy of the weather, diseases, and uncertainty about the prices they'll receive for their products. There are very few businesses (besides such high-risk/high-payoff industries as gold mining and oil drilling) in which managers are willing to invest money and resources at the start of a production cycle without some clear indication of the likely consequences of their decisions. Farmers have historically been "price takers." However, as farm businesses become fewer and larger, surviving farm operators are using risk-reducing strategies like crop insurance and forward contracting to market hogs to reduce some of the uncertainty in the planning process.

Despite our best efforts to control swine production, however, there is an element of chance in the outcome of almost everything we do. *Decision analysis* is a formal, structured way to model those chance events. It uses a *decision tree* as a pictorial representation of the flow of events in a logical and time-sequenced manner, so that the decision maker can consider the probabilities of each outcome. It quantifies and helps us consider the effects of chance on the outcome of a given decision.

In using decision analysis, it is important to understand that the objective is not to make a prediction about the fate of an individual animal. A sick sow may still die despite being treated, or may live without treatment. Decision analysis uses probabilities and monetary values to provide a guide for what should be done. If the values you assign to the tree accurately reflect the real world, you can use the same generic decision tree in many identical situations which, regardless of the outcome of any single decision, will guide you to recommend the more profitable options over the long run.

Department of Clinical & Population Sciences, University of Minnesota, College of Veterinary Medicine, St. Paul, Minnesota 55108.

Using decision analysis to advise clients

What is the problem?

The first step in building a decision tree is to define the problem. Fig 1 is a decision tree of a problem familiar to all veterinarians: what to do when confronted with a sick animal. Suppose that a sow is diagnosed as diseased. If not treated, there is a 40% chance that she will live. There is a treatment available that costs \$50, and raises the **probability** of survival to 80%. The owner values the sow at \$300 if she lives. If she dies, it will cost \$20 to dispose of the carcass. To treat, or not to treat? That is the question.

Probability: numerical values indicating the likelihood of a given outcome. Often given as a number in the range 0 to 1, where 0.30 would indicate a 30% or a "30 out of 100" chance.

What can we do?

The next step in building a decision tree is to identify a mutually exclusive, exhaustive list of all possible courses of action to address the problem. In the case of a sick sow, youcan recommend that the sowbe:

- treated;
- left untreated;
- immediately culled; or

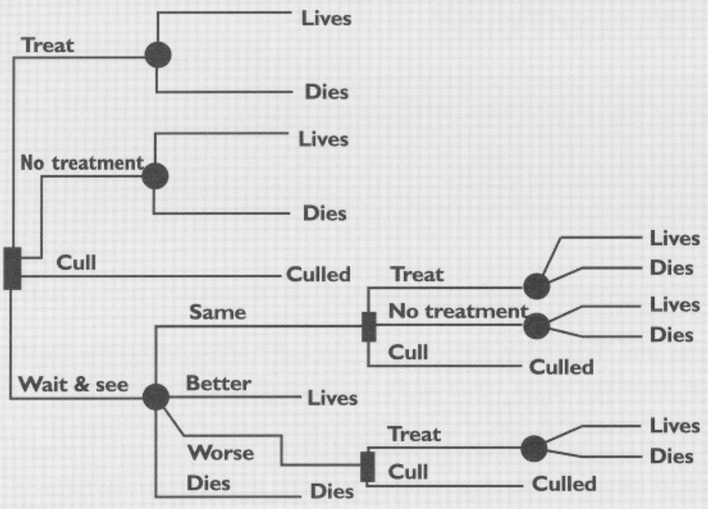

Fig 1.—A decision tree showing oprtions for a diseased sow.

- that the producer wait and see, deferring the decision until later.

The first node in a decision tree, then, is always a decision node, which is conventionally represented as a rectangular box (Fig 1). A separate branch (represented as lines) emanates from the decision node (to the right) for each possible decision that is under consideration. Thus, the model is exhaustive in that it includes all possible choices to be considered in the decision. The model is also exclusive in the sense that only one choice may be selected at each decision node.

Is another decision implicit in a given decision node?

Sometimes, branches emanating from a decision node can lead to other decision nodes. For example, if the decision maker "waits and sees," the sow may:

- stay the same;
- get better;
- get worse; or
- die.

In this model, if the sow gets better, no further decision is needed. If however, she stays the same or gets worse under observation, the veterinarian must decide whether to treat it, leave it untreated, or cull it. Depending upon the complexity of a given problem, there can be decision nodes at many stages of the decision tree. Part of the challenge of using decision trees is to adequately delimit the initial problem so that the resulting decision tree is not too cumbersome to use. In this example, we have deliberately trimmed the "No treatment" branch as a consequence of waiting and the animal getting worse, because it is not normally considered a feasible action.

How likely is any particular outcome?

Once all decision nodes are in place, you must identify an exhaustive and mutually exclusive list of all the possible outcomes of each decision. Each of these possible outcomes will branch out of a chance node (represented as a circle). Branches to the right of a chance node should represent all possible outcomes that can result at that point. Each outcome (i.e. each branch stemming from a chance node) has an associated probability, indicating the likelihood that the particular outcome represented by that branch will occur. In our example in Fig 1, all the possible outcomes of three of our four original options are shown emanating from their respective chance nodes. It is important to realize (and to emphasize to your clients) that once you move from a decision node to a given chance node (i.e. once you choose a course of action), the outcome is beyond the control of the decision maker. Note that in the case of the fourth option

Features of decision analysis

Decision analysis is explicit because it forces the decision maker to separate the problem into its component parts, without losing the context of the "big picture." The analytic approach forces the decision maker to consider explicitly the timing of choices that must be made, and the data that must be acquired to make informed decisions. Uncertainties involved and the relative values of possible outcomes are also explicitly stated.

Decision analysis is quantitative because the decision maker is forced to use the language of probability. Probabilities are used to express a strength of belief that an animal will live or die, or whether a herd will experience a clinical outbreak of disease. Anecdotal estimates such as "very likely," or "a good chance that...," are replaced by numerical estimates of probabilities such as "an 80% chance of success." This approach is particularly helpful in communicating the opinion of the veterinarian to the animal owner, who is usually the ultimate decision maker acting under the influence of veterinary advice.

Decision analysis is prescriptive, rather than descriptive. It is intended to aid the decision maker in deciding what *should* be done under a given set of circumstances. The decision is consistent with the problem as it has been laid out logically as a decision tree. The uncertainties involved have been identified and quantified, and relative values placed upon possible outcomes.

(Cull), there is no chance element. In this model, we assume that if the decision is made to cull, the sow will live to be slaughtered.

How do I quantify the likelihood of a given outcome?

If, as in the case of an acute disease, the producer is unwilling to consider culling or waiting to see how the condition develops, it may be appropriate to trim the decision tree down to only two branches at the first decision node. If you choose to treat the animal, there is a probability (P1) that it will live ("Lives"), and a corresponding probability (1–P1) that it will die ("Dies") (Fig 2). Probabilities are quantified by a value ranging from zero to one. A probability of zero indicates that the event will never happen, while a probability of one indicates that the event will certainly occur. A probability of 0.3 indicates that the outcome is expected to occur 30% of the time. The sum of the probabilities of all outcome branches emanating from a single chance node must equal one. You will need to assign a probability value to each outcome branch based on your experience and knowledge as a swine practitioner. At this stage of decision analysis, this can be a reasonable "ballpark" guess, since you will have the opportunity later to assess the precision of your estimates. You can also use conclusions from current literature on a subject or the data from production record systems to determine probability values for some decision-making tasks.

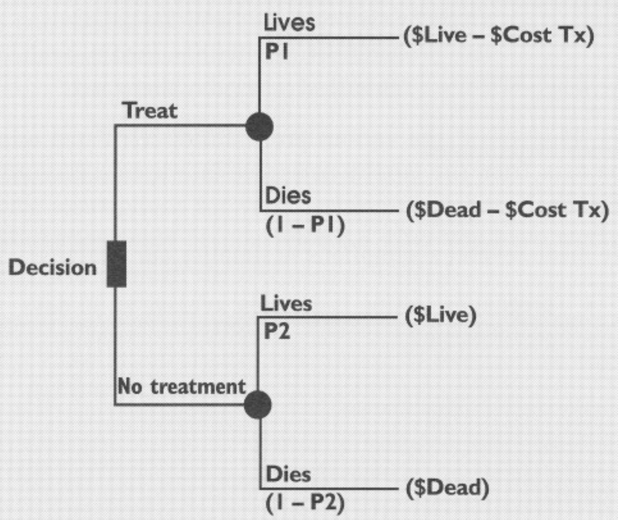

Fig 2.— A trimmed decision tree showing probability and value variables for each outcome.

How do I select the best choice of action?

All outcome branches end in terminal nodes, which represent the consequences of each decision. Each terminal node is assigned a representation of the expected monetary value of each decision: either a dollar value or an equation of several separate but component variables. It is important to include all the pertinent financial components of each terminal node to get an accurate sense of the ultimate cost or gain for each decision. In our example, we have expressed expected monetary values as combinations of variables: the value of a live animal ($Live), the value (often negative) of a dead animal ($Dead), and the cost of treatment ($Cost Tx.). Actual probabilities and terminal node dollar values, appropriate for a given farm, can be "plugged in" to these equations to customize the decision tree so that it can be used on any farm. These values will vary from farm to farm depending upon the severity of the disease (case fatality rate), the efficacy and cost of treatment(s), and the relative values placed upon live and dead animals.

> **Folding back:** multiplying associated probabilities by the monetary values at each terminal node.
>
> **Expected value:** sum of all values that result from folding back a branch leading out of a decision node.

Once you have substituted actual numerical values for the variables in the tree, you can evaluate the possible decisions by a process called **folding back,** accomplished by multiplying associated probabilities by the monetary values at each terminal node. All branches emanating from the same node are folded back and the results summed (Fig 3, Calculation 1). This sum is called th**e expected value** of that node. In a larger tree, you must fold back until you have reached the original decision node. You can use the expected values that result from the folding back process as the basis to decide the most desirable course of action.

Decision analysis allows you to compare the financial worth of all possible outcomes. Because the expected value of treating an animal is $78 more than the expected value of not treating it ($186–$108 = $78), you would advise your client, based on decision analysis, to treat the animal (Fig 3).

Developing a model

How confident can I be that I have assigned realistic values to the model and that the outcome I advocate is actually the best choice?

Let's use another example scenario — whether or not to use an ultrasound test for pregnancy in a herd — to assess how much we can rely on the results of our decision analysis tree. This will help us ascertain how the situation in the herd would have to change before we would want to change our advice to the client.

In many swine herds, exposing bred sows and gilts to a boar every day is still the preferred method of pregnancy detection. Under this system, animals that fail to show signs of estrus after they are bred are presumed pregnant. With this method, however, false positive results often occur if non-

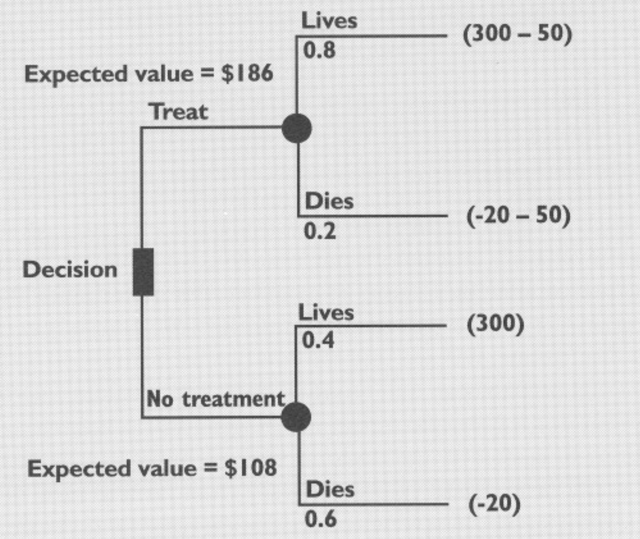

Fig 3.— A decision tree with values assigned.

Expected value (treated and lives) =	0.8 × (300 – 50) =	200
Expected value (treated and dies) =	0.2 × (-20 – 50) =	-14
Expected value of treating =		186
Expected value (not treated and lives) =	0.4 × 300 =	120
Expected value (not treated and dies) =	0.6 × -20 =	-12
Expected value of not treating =		108

Calculation 1.— Calculating example of expected values.

pregnant females fail to show strong signs of estrus, return to estrus at irregular intervals, or become anestrous. Mechanical ultrasound devices and hormone assessment kits have been developed to improve the accuracy of pregnancy diagnosis in bred sows and gilts, and are available to veterinarians and producers.

Unfortunately, these diagnostic tests are still not perfect. There is still the risk that a pregnant female will be classified as not pregnant (false negative), or a that nonpregnant animal will be classified as pregnant (false positive). Someone deciding whether to use one of these tests must know the likelihood that "positive" sows are actually pregnant and that "negative" sows are truly not pregnant. **Predictive values** quantify these probabilities.

The predictive values of a test depend upon the **specificity, sensitivity**, and **prevalence.** Sensitivity and specificity are innate characteristics of a test and never vary. However, the prevalence of disease in the population will affect the proportion of test-positive sows and gilts that are actually diseased. Similarly, the probabilities of obtaining a positive or negative test will also change. In our example the "disease" of interest is pregnancy.

In order to calculate probabilities necessary to construct an ultrasound testing decision tree, we'll use data Almond and

Table 1.— Calculation of probabilities.

	Actual status		
Test result	+	-	Total
+	38	4	42
-	2	6	8
Total	40	10	50

Predictive value positive =	38 ÷ 42 =	90.5%
Predictive value negative =	6 ÷ 8 =	75.0%
Pregnancy rate =	40 ÷ 50 =	80.0%
Probability of positive test =	42 ÷ 50 =	84.0%
Probability of negative test =	8 ÷ 50 =	16.0%

Dial[1] published after comparing seven pregnancy-detection procedures. Each procedure was applied to groups of 50 sows by three independent experts. The results for each procedure were compared to the number of sows that subsequently farrowed consistent with being mated 31–35 days prior to being tested for pregnancy (i.e. the "prevalence" of pregnancy in this herd). We summarized their data for the oscilloscope amplitude-depth ultrasound pregnancy detector in a two-by-two table to facilitate calculation of five probabilities that will be used to evaluate the diagnostic test (Table 1).

We "plugged in" the various probabilities calculated using the data from the two-by-two table into the decision tree (Fig 4).

We made several assumptions to simplify our analysis:

- pregnancy testing is done on average at 30 days post breeding;

Predictive values: numerical values given to the probability that a test is correct.

Specificity: (True negative rate.) The proportion of animals identified by the test as diseased that are truly nondiseased. Low specificity results in false positives.

Sensitivity: (True positive rate.) The proportion of animals identified by the test as diseased that are truly diseased. Low sensitivity results in false negatives.

Prevalence (of a disease or condition): The proportion of animals in a given population with a specific disease or condition at a specific time.

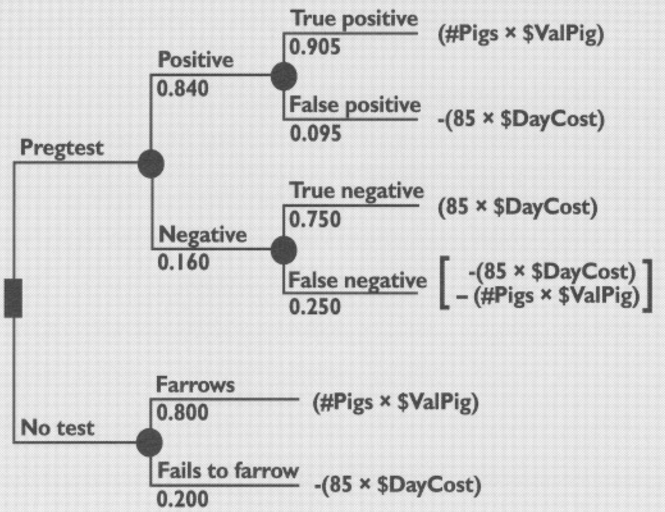

Fig 4.— Decision tree with probabilities "plugged in."

- all sows with positive test results will be kept until their due-to-farrow dates, unless they die or are culled for nonreproductive reasons;
- all sows with negative test results will be culled immediately; and
- income from the sale of a cull sow ($125) exactly equals the additional cost of retaining a replacement gilt in the herd that would otherwise be sold as a market hog.

(You will also probably need to make some assumptions for your decision trees, but often these assumptions only become apparent after you have worked your way through the tree once or twice.)

At the decision node, you'll need to decide whether to test or not to test. If you decide not to test, then we will assume sows will remain in the herd until their due-to-farrow dates. Then, the proportion of sows expected to be pregnant (80%) will farrow, while 20% will fail to farrow. The expected value of a farrowing sow is the average weaned litter size for the herd multiplied by the value of a pig. In reality, not all sows testing positive will remain in the herd until farrowing. Some may die, and others may be culled for lameness or other nonreproductive reasons. Because these losses will occur regardless of whether or not pregnancy is tested, we could safely ignore them for the purpose of evaluating the financial benefits of testing. However, in the examples that follow, we will use the adjusted farrowing rate (AFR) as the best estimate of a herd's pregnancy rate.

As was the case for our first example, some of the chance nodes can lead to other chance nodes in this decision tree. For example, once the decision to use the test on a sow is made, the test will yield either a positive or a negative result. That result may be either true or false. The monetary values at the terminal nodes of the branches emanating from the positive test are similar to those for the "no-test" scenario. This is logically consistent, as sows that test positive are also assumed to stay in the herd until their due-to-farrow dates.

Adjusted farrowing rate (AFR) is calculated by the PigCHAMP® program as:

$$\frac{\text{No. of sows farrowed}}{\text{(No. gilts \& sows served – No. removed for non-reproductive reasons)}}$$

...where the period in which the farrowings occur is compared to a period of equal length in which the services occurred 115 days earlier.

Because a negative test result dictates an instant cull, the monetary values associated with a negative test include a savings of 85 days' variable costs that would otherwise be incurred by keeping the sow until her due-to-farrow date. However, the additional penalty incurred by culling a pregnant sow based on a false-negative test result is the value of the litter that has been lost.

Now that the logical structure is complete, and the probabilities assigned, the next step is to assign values to the variables shown in expected value expressions. These will

"How much should the owner of the sow be prepared to pay for a treatment?"

To answer this question, you must perform a special type of sensitivity analysis to your decision tree, called "threshold analysis." Threshold analysis is performed by repeatedly substituting different values (by trial and error) for the cost of treatment until the expected values of each outcome are exactly equal. You can use a specialized computer program, such as SMLTREE (Hollenberg J. *Smltree – The All-Purpose Decision Tree Builder.* [Computer Software and Manual] 445 E 68th St., Box 20, New York, NY 10021), or an electronic spreadsheet program to set up your tree to perform this analysis for you. In our first example, the threshold of treatment cost (Cost Tx) is $108. If this were the case, the expected values of treating or not treating would both be $108 and, in theory, the owner would be unable to choose between the two alternatives. You may find it worthwhile to spend some time experimenting with various combinations of costs and efficacies of alternative treatments to revise your own scale of charges!

Table 2.— Corrected probabilities for decision analysis model and break-even cost per sow of using pregnancy testing at various adjusted farrowing rates. Litter valued at $70; DayCost = $1.00.

	Adjusted farrowing rate (%)									
	50	55	60	65	70	75	80	85	90	95
Predictive value positive	0.704	0.744	0.781	0.815	0.847	0.877	0.905	0.931	0.955	0.978
Predictive value negative	0.923	0.908	0.889	0.866	0.837	0.800	0.750	0.679	0.571	0.385
Probability of positive test	0.675	0.703	0.730	0.758	0.785	0.813	0.840	0.868	0.895	0.922
Expected value of testing	$37.90	$40.00	$42.08	$44.11	$46.20	$48.33	$50.43	$52.53	$54.52	$56.51
Expected value of not testing	(-$7.50)	$0.25	$8.00	$15.75	$23.50	$31.25	$39.00	$46.75	$54.50	$62.25
Break-even cost of testing	$45.40	$39.75	$34.08	$28.36	$22.70	$17.08	$11.43	$5.78	$0.02	(-$9.99)

vary from herd to herd. Again to simplify the analysis, a few more assumptions are made:

- Sow feed = 6 lb (2.72 kg) per day @ $160 per ton = 48¢ per day. Nonfeed costs = 52¢ per day. Therefore, the cost of keeping a sow in the herd for one day ("$DayCost") = $1.00.
- Records for this farm show an average liveborn litter size of 8.05, and an average weaned litter size of 7.0. As the output of the breeding herd is measured in weaned pigs, we will count only those liveborn pigs that will live to be weaned. Therefore, "#Pigs" = 7.0.
- Placing a value on a baby pig is a contentious issue among economists. My personal belief is that an appropriate value for a growing animal is the margin that would be realized by raising and selling it. If I could realize a profit of $20 on raising a pig, then I would be willing to pay up to, but not more than, $20 to purchase it. (Conversely, if I were losing money in the hog business, every live pig produced would represent a liability to my business, rather than an asset.) In this example, "$ValPig" = $10.00.

Folding back the tree using the values noted above, we have:

- expected value of testing = $50.43; versus
- expected value of not testing = $39.00

Because the expected value of ultrasonic pregnancy testing is greater than not testing, it is the preferred choice in this case. Further, we can say that (in this scenario our) analysis shows that the benefit of testing is $11.43 ($50.43 – $39.00) per sow tested. As a charge for the ultrasound test was not included in the analysis, we can say (and tell our clients if we are providing this service), that it is worthwhile to pay up to $11.43 per test to use the procedure in this herd, assuming an adjusted farrowing rate of 80%. The difference in the expected values is the *maximum* amount that the producer should be willing to pay for the test. If the test is cheap ($3), the producer will receive a four-fold return on this investment. If the test costs $10, however, the margin would be slim and the producer may wish to invest the money elsewhere!

You probably would like to know how confident you could be in advocating that this producer use ultrasound to pregnancy-test the sows. The producer may also want to know whether to continue to use pregnancy testing if hog prices drop significantly. You can perform a sensitivity analysis to determine how confident you can be in your advice, and how the situation in the herd would need to change before you would change your advice. We did a sensitivity analysis by varying AFR while holding $DayCost, #Pigs and $ValPig constant (Table 2). Because the predictive values and probability of a positive test change with prevalence, corrected probabilities have been calculated for AFRs between 50% and 95%. Our test was evaluated in a herd with an 80% AFR (Table 2 — highlighted column). As farrowing rate increases, the predictive value of a positive test result increases, while the predictive value of a negative test result falls. Also, the higher the adjusted farrowing rate, the greater the probability of a positive test result.

In the extreme case, where the adjusted farrowing rate is only 50%, the predictive value of a negative test result rises to 92.3%, while the predictive value of a positive test result is only 70.4%. That is, out of every 100 positive test results, only 70 sows would be actually pregnant, while 30 others would be false positives, using this test. Even though the proportion of false positive results increases as farrowing rates fall, the net financial benefit of testing also increases. This makes intuitive sense. If the adjusted farrowing rate were 100%, even a perfect test would have no value to a decision maker, because the outcome would be certain. Inspection of Table 1 shows that under our assumptions, the herd manager should pay up to $5.78 per sow to use this test at an AFR of 85%, but only 2¢ at an AFR of 90%.

The next logical step is to perform sensitivity analyses changing the values of $DayCost, $ValPig, and #Pigs to determine how the scenario might have to change before it would be best to decide not to use the test (Table 3). To simplify the table, $ValPig and #Pigs have been combined as a single variable representing the dollar value of the litter of pigs. As the value of the litter increases in relation to $DayCost, the break-even cost of routine ultrasonic pregnancy testing is reduced. If $DayCost is set at $1.00, then the break-even cost is equal to $1.56 at an AFR of 85% and a litter value of $120. However, when $DayCost is set at $1.50 (in the same herd) the owner could afford to pay up to $7.40 per test where the average litter is valued at $120.

Routine pregnancy testing may not be financially justified in herds with very high farrowing rates and low housing and feed costs. The higher the value placed upon the litter of pigs in relation to the costs of keeping and replacing nonpregnant sows and gilts, the lower the break-even cost of testing. However, as the actual cost of pregnancy testing is relatively low (labor plus depreciation and maintenance of the machine) it would appear to be a cost-effective activity in most situations where the AFR is 75% or less (Table 3), given the data describing the characteristics of this particular test published in the literature. Generally, the lower the adjusted farrowing rate of the herd, and the higher the cost of keeping a nonpregnant sow or gilt in the herd, the greater the potential benefit of routine testing.

The analyses presented in this paper are deliberately simple because my intention is to present the basic principles of decision tree analysis. Once you understand the basic principles, it is a relatively simple process to add more branches to the pregnancy diagnosis model to deal with the probability of a sow returning to estrus, or to expand the utility equations to account for differences in values of a culled sow and a purchased replacement gilt.

Once the logical structure of any decision tree is set up, it can be applied to different herd situations by "plugging in "farm- and test-specific data to tailor it closely to real-world situations. Most of the data required to complete and analyze the problem may already exist in the farm's record system, or can be estimated with a little creative thought and effort. The graphic layout and the explicit nature of the language of probability can be used to reduce complex decisions to simple terms. Every day, veterinarians make decisions using their best judgement. Decision analysis can help formalize the decision-making process by replacing emotion or 'gut feeling' with logic. It also offers an excellent means of communicating with and involving clients in the decision-making process.

Reference

1. Almond GW, and Dial GD. Pregnancy diagnosis in swine: A comparison of the accuracies of mechanical and endocrine tests with return to estrus. *JAVMA* 1986; 189:1567-71.

Table 3.— Sensitivity analysis showing the break-even cost per sow of using ultrasonic pregnancy testing at various litter values and two values of DayCost.

Value of litter	Break-even cost at DayCost = $1.00					Break-even cost at DayCost = $1.50				
	Adjusted farrowing rate (%)									
	70	75	80	85	90	70	75	80	85	90
$200	$13.58	7.41	1.05	-	-	$27.38	18.55	9.75	0.66	-
180	14.98	8.90	2.65	-	-	28.79	20.03	11.16	2.34	-
160	16.38	10.38	4.25	-	-	30.19	21.52	12.76	1.03	-
140	17.79	11.87	5.85	-	-	31.59	23.00	14.35	5.72	-
120	19.19	13.36	7.44	1.56	-	32.99	24.50	15.95	7.40	-
100	20.59	14.85	9.04	3.25	-	34.40	25.99	17.55	9.08	0.48
80	22.00	16.34	10.63	4.93	-	35.80	27.47	19.14	10.77	2.29
60	23.40	17.82	12.23	6.62	0.92	37.20	28.96	20.74	12.45	4.10
40	24.80	19.31	13.83	8.31	2.73	38.60	30.46	22.33	14.14	5.90
20	26.21	20.80	15.42	9.89	4.54	40.00	31.95	23.93	15.82	7.71