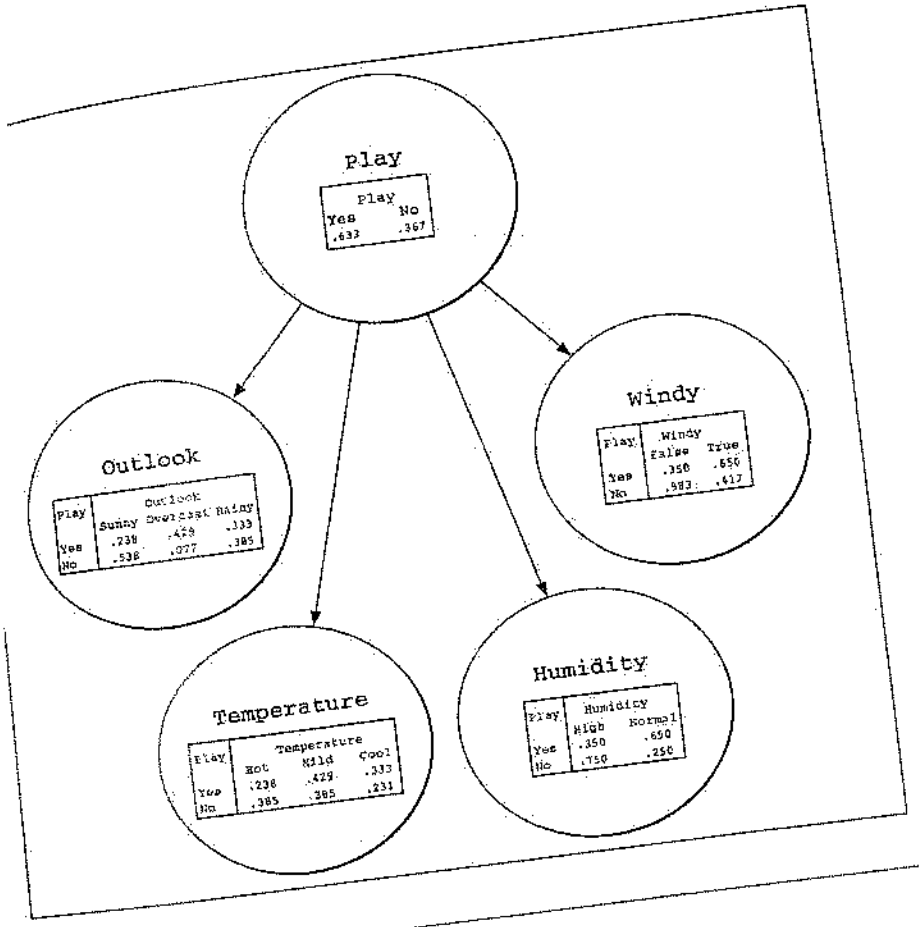

FIGURE 9.1
A simple Bayesian network for the weather data.

other nodes. But in Bayesian networks the structure of the graph is only half the story. Fig. 9.1 shows a table inside each node. The information in the tables defines a probability distribution that is used to predict the class probabilities for any given instance.

Before looking at how to compute this probability distribution, consider the information in the tables. The lower four tables (for *outlook*, *temperature*, *humidity*, and *windy*) have two parts separated by a vertical line. On the left are the values of *play*, and on the right are the corresponding probabilities for each value of the attribute represented by the node. In general, the left side contains a column for every edge pointing to the node; in this case just an edge emanating from the node for the *play* attribute. That is why the table associated with *play* itself does not have a left side: it has no parents. In general, each row of probabilities corresponds to one combination of values of the parent attributes, and the entries in the row show the probability of each value of the node's attribute given this

combination. In effect, each row defines a probability distribution over the values of the node's attribute. The entries in a row always sum to 1.

Fig. 9.2 shows a more complex network for the same problem, where three nodes (*windy*, *temperature*, and *humidity*) have two parents. Again, there is one column on the left for each parent and as many columns on the right as the attribute has values. Consider the first row of the table associated with the

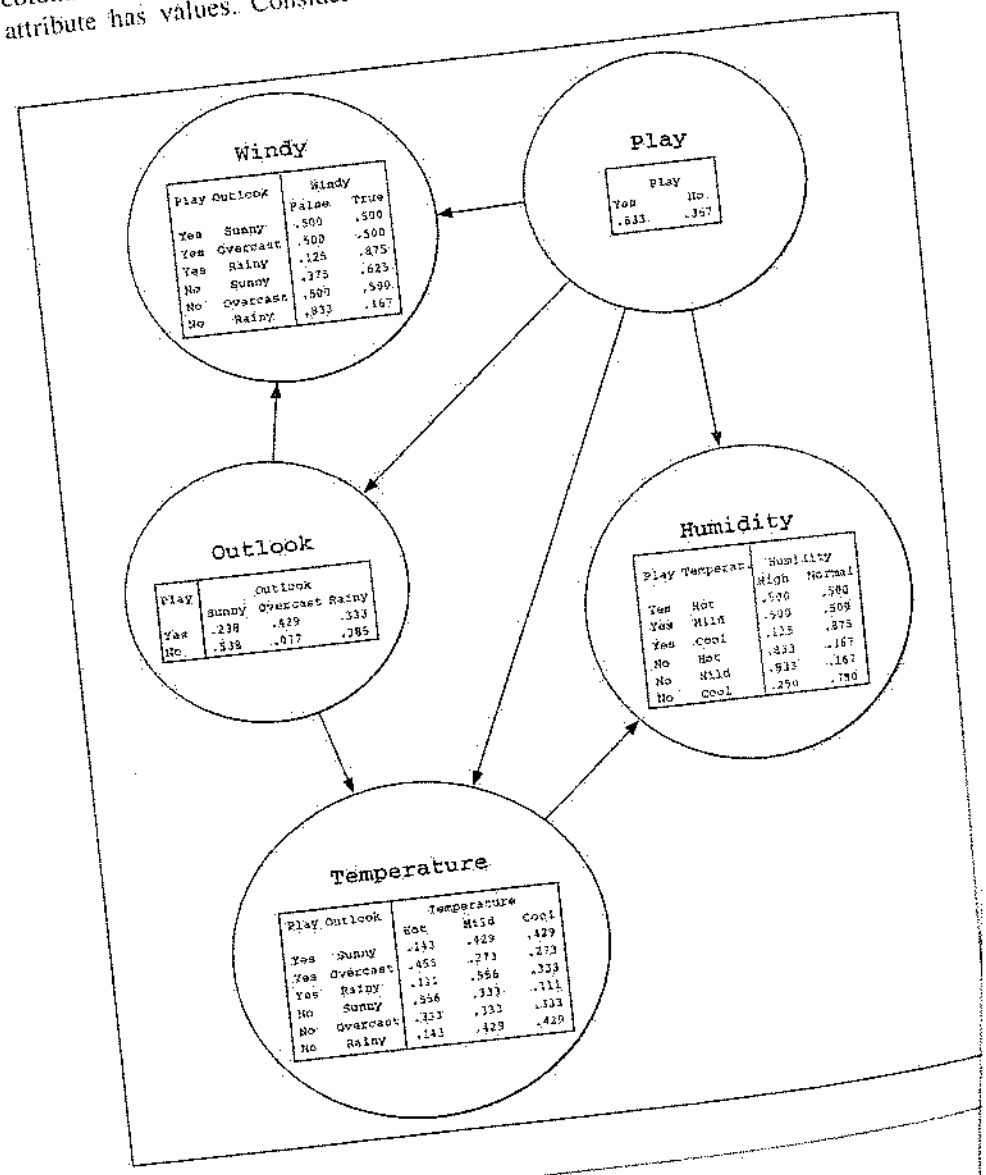

FIGURE 9.2
Another Bayesian network for the weather data.