

<i>Cold</i>	<i>Flu</i>	<i>Malaria</i>	$P(\text{fever} \cdot)$	$P(\neg\text{fever} \cdot)$
<i>f</i>	<i>f</i>	<i>f</i>	0.0	1.0
<i>f</i>	<i>f</i>	<i>t</i>	0.9	<b>0.1</b>
<i>f</i>	<i>t</i>	<i>f</i>	0.8	<b>0.2</b>
<i>f</i>	<i>t</i>	<i>t</i>	0.98	$0.02 = 0.2 \times 0.1$
<i>t</i>	<i>f</i>	<i>f</i>	0.4	<b>0.6</b>
<i>t</i>	<i>f</i>	<i>t</i>	0.94	$0.06 = 0.6 \times 0.1$
<i>t</i>	<i>t</i>	<i>f</i>	0.88	$0.12 = 0.6 \times 0.2$
<i>t</i>	<i>t</i>	<i>t</i>	0.988	$0.012 = 0.6 \times 0.2 \times 0.1$

**Figure 13.5** A complete conditional probability table for  $\mathbf{P}(\text{Fever} | \text{Cold}, \text{Flu}, \text{Malaria})$ , assuming a noisy-OR model with the three  $q$ -values shown in bold.

are true. The noisy-OR model allows for uncertainty about the ability of each parent to cause the child to be true—the causal relationship between parent and child may be *inhibited*, and so a patient could have a cold, but not exhibit a fever.

The model makes two assumptions. First, it assumes that all the possible causes are listed. (If some are missing, we can always add a so-called **leak node** that covers “miscellaneous causes.”) Second, it assumes that inhibition of each parent is independent of inhibition of any other parents: for example, whatever inhibits *Malaria* from causing a fever is independent of whatever inhibits *Flu* from causing a fever. Given these assumptions, *Fever* is *false* if and only if all its *true* parents are inhibited, and the probability of this is the product of the inhibition probabilities  $q_j$  for each parent. Let us suppose these individual inhibition probabilities are as follows:

$$q_{\text{cold}} = P(\neg\text{fever} | \text{cold}, \neg\text{flu}, \neg\text{malaria}) = 0.6,$$

$$q_{\text{flu}} = P(\neg\text{fever} | \neg\text{cold}, \text{flu}, \neg\text{malaria}) = 0.2,$$

$$q_{\text{malaria}} = P(\neg\text{fever} | \neg\text{cold}, \neg\text{flu}, \text{malaria}) = 0.1.$$

Then, from this information and the noisy-OR assumptions, the entire CPT can be built. The general rule is that

$$P(x_i | \text{parents}(X_i)) = 1 - \prod_{\{j: X_j = \text{true}\}} q_j,$$

where the product is taken over the parents that are set to true for that row of the CPT. Figure 13.5 illustrates this calculation.

In general, noisy logical relationships in which a variable depends on  $k$  parents can be described using  $O(k)$  parameters instead of  $O(2^k)$  for the full conditional probability table. This makes assessment and learning much easier. For example, the CPCS network (Pradhan *et al.*, 1994) uses noisy-OR and noisy-MAX distributions to model relationships among diseases and symptoms in internal medicine. With 448 nodes and 906 links, it requires only 8,254 parameters instead of 133,931,430 for a network with full CPTs.