Fundamentals of Artificial Intelligence

Laboratory

Dr. Mauro Dragoni

Department of Information Engineering and Computer Science

Academic Year 2020/2021

Exercise 13.1

- Given the full joint distribution shown in the figure below, calculate the following:

  a) **P**(toothache)
  b) **P**(Cavity)
  c) **P**(Toothache | cavity)
  d) **P**(Cavity | toothache ∨ catch)

	toothache		*¬toothache*	
	catch	*¬catch*	*catch*	*¬catch*
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Figure 13.3 A full joint distribution for the *Toothache*, *Cavity*, *Catch* world.

Exercise 13.1

The main point of this exercise is to understand the various notations of bold versus non-bold P, and uppercase versus lowercase variable names. The rest is easy, involving a small matter of addition.

a. This asks for the probability that Toothache is true.
   P(toothache) = 0.108 + 0.012 + 0.016 + 0.064 = 0.2

b. This asks for the vector of probability values for the random variable Cavity.
   It has two values, which we list in the order ⟨true, false⟩.
   First add up 0.108 + 0.012 + 0.072 + 0.008 = 0.2. Then we have P(Cavity) = ⟨0.2, 0.8⟩ .

c. This asks for the vector of probability values for Toothache, given that Cavity is true.
   P(Toothache | cavity) = ⟨(0.108 + 0.012) / 0.2, (0.072 + 0.008) / 0.2⟩ = ⟨0.6, 0.4⟩

d. This asks for the vector of probability values for Cavity, given that either Toothache or Catch is true.
   First compute P(toothache ∨ catch) = 0.108 + 0.012 + 0.016 + 0.064 + 0.072 + 0.144 = 0.416.
   Then P(Cavity | toothache ∨ catch) = ⟨(0.108 + 0.012 + 0.072) / 0.416, (0.016 + 0.064 + 0.144) / 0.416⟩
   = ⟨0.4615, 0.5384⟩

Exercise 13.2

- Consider two medical tests, A and B, for a virus.
- Test A is 95% effective at recognizing the virus when it is present, but has a 10% false positive rate (indicating that the virus is present, when it is not).
- Test B is 90% effective at recognizing the virus, but has a 5% false positive rate.
- The two tests use independent methods of identifying the virus.
- The virus is carried by 1% of all people.
- Say that a person is tested for the virus using only one of the tests, and that test comes back positive for carrying the virus.
- Which test returning positive is more indicative of someone really carrying the virus? Justify your answer mathematically.

Exercise 13.2

Let V be the statement that the patient has the virus, and A and B the statements that the medical tests A and B returned positive, respectively.

The problem statement gives:

$P(V) = 0.01$
$P(A|V) = 0.95$
$P(A|\neg V) = 0.10$
$P(B|V) = 0.90$
$P(B|\neg V) = 0.05$

The test whose positive result is more indicative of the virus being present is the one whose posterior probability, $P(V|A)$ or $P(V|B)$ is largest.

One can compute these probabilities directly from the information given, finding that $P(V|A) = 0.0876$ and $P(V|B) = 0.1538$, so B is more indicative.

Exercise 13.3

- Consider the following Bayesian Network, where variables **A-D** are all Boolean-valued:

P(A=true) = 0.7

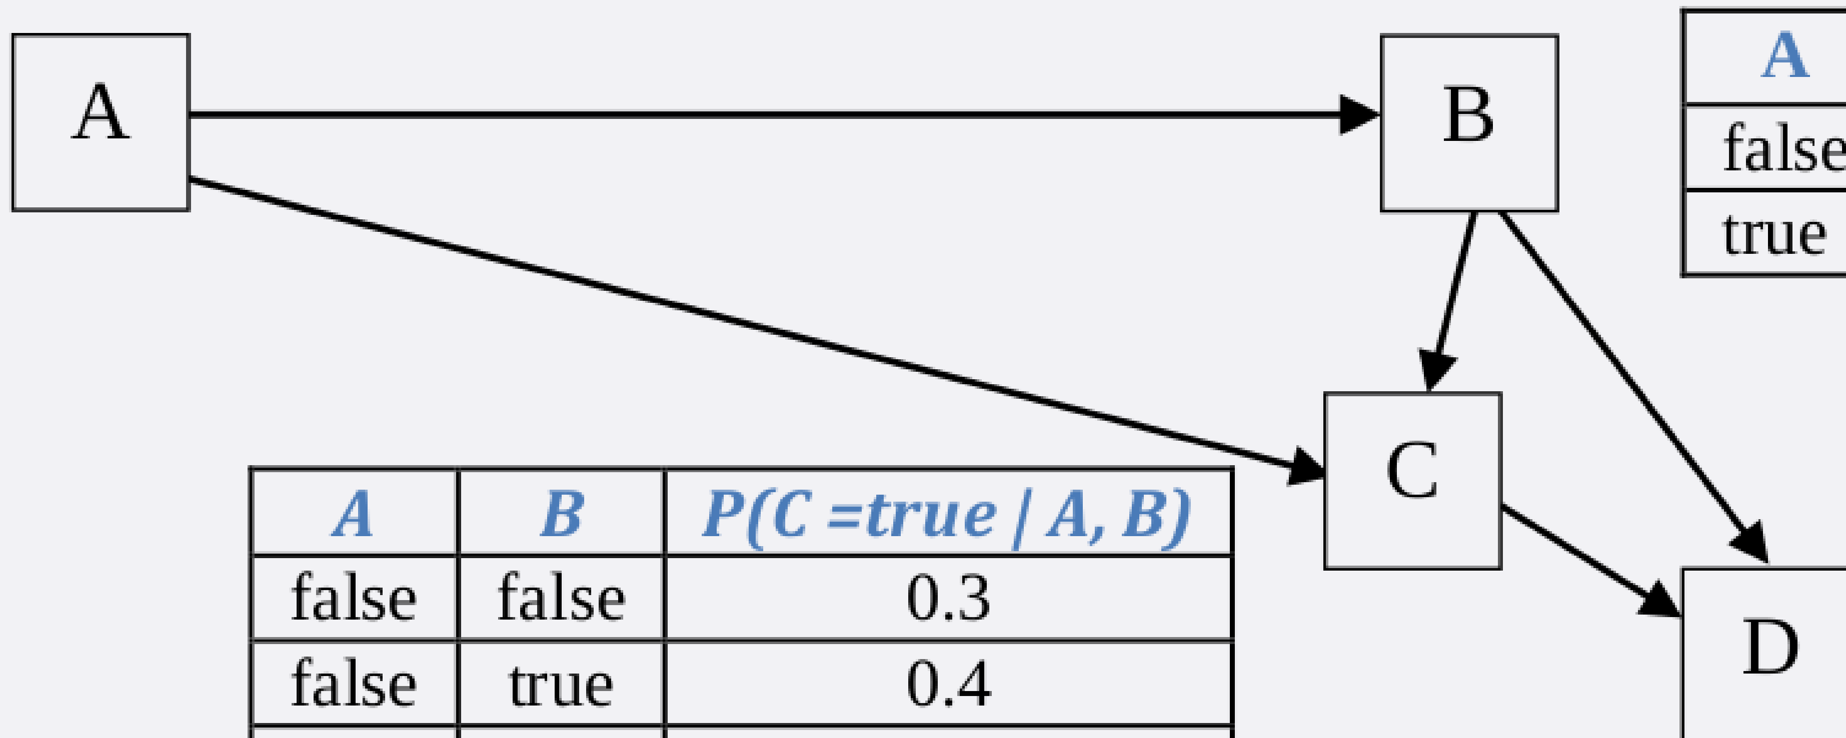

A	P(B =true | A)
false	0.1
true	0.8

A	*B*	*P(C =true | A, B)*
false	false	0.3
false	true	0.4
true	false	0.5
true	true	0.6

B	*C*	*P(D =true | B,C)*
false	false	0.9
false	true	0.3
true	false	0.4
true	true	0.2

- Show your work for the following calculations.
  1. Compute *P*(A=true and B=true and C=true and D=true).
  2. Compute *P*(B=false and C=false and D=false).
  3. Compute *P*(A=true | B=true and C=true and D=false).
  4. Compute *P*(D=false | A=true and B=true and C=true).
     (Do this one without employing the Markov Blanket property, i.e., algebraically or numerically confirm that knowing A=true does not change the results from computing *P*(D=false | B=true and C=true).
  5. Compute *P*((A=true and D=true) or (B=true and C=true)).

Exercise 13.3

P(A=true) = 0.7

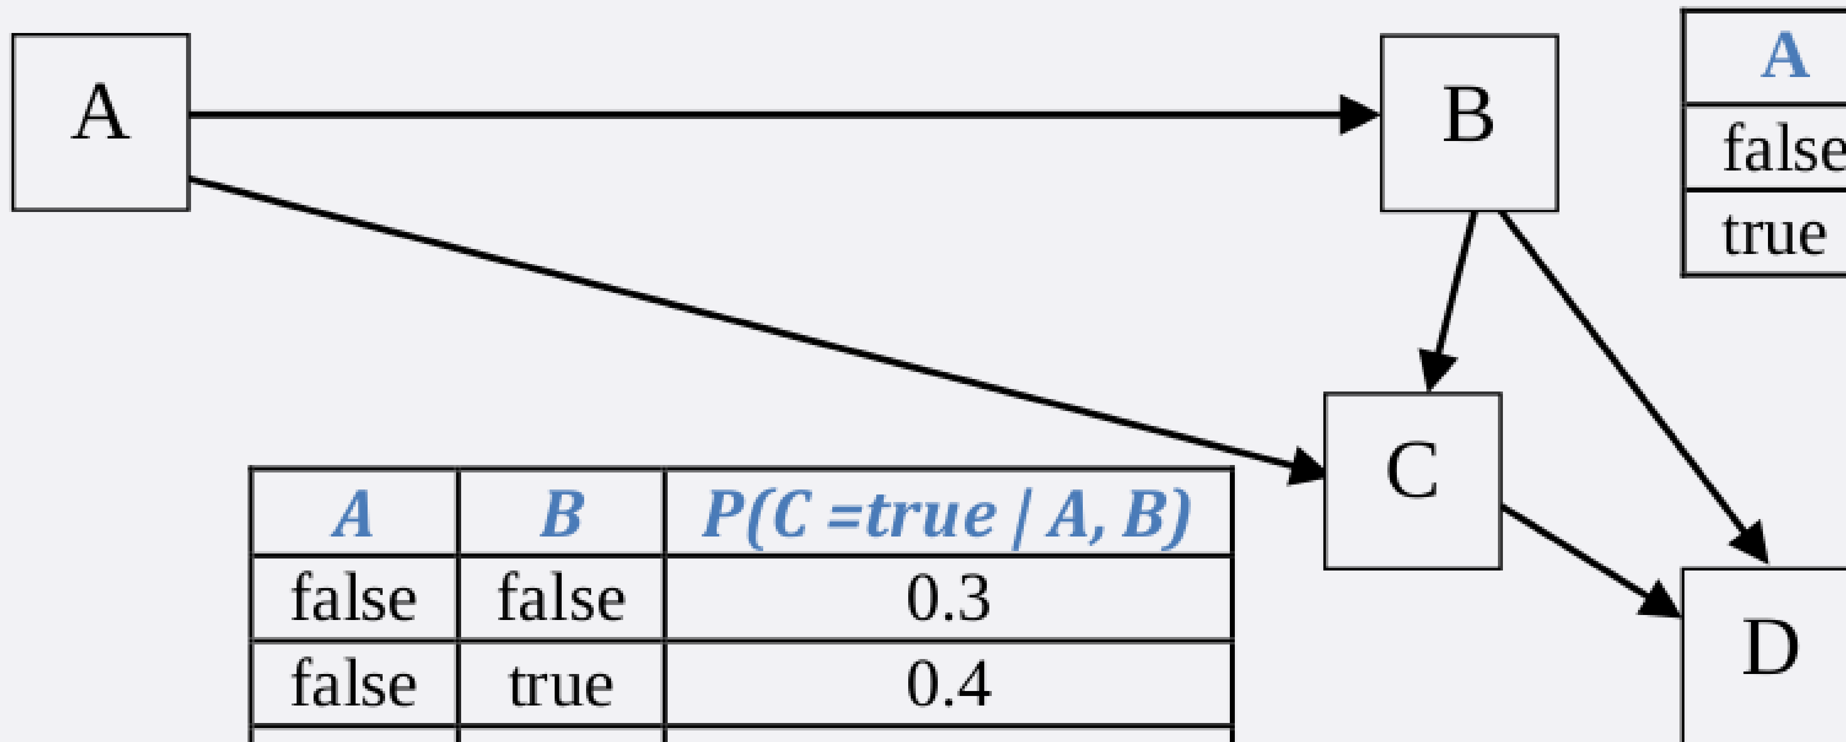

A	P(B =true | A)
false	0.1
true	0.8

A	*B*	*P(C =true | A, B)*
false	false	0.3
false	true	0.4
true	false	0.5
true	true	0.6

B	*C*	*P(D =true | B,C)*
false	false	0.9
false	true	0.3
true	false	0.4
true	true	0.2

1. Compute *P*(A=true and B=true and C=true and D=true).

P(A,B,C,D) = P(A) P(B|A) P(C|A,B) P(D|B,C) = 0.7 * 0.8 * 0.6 * 0.2 = 0.0672

Exercise 13.3

P(A=true) = 0.7

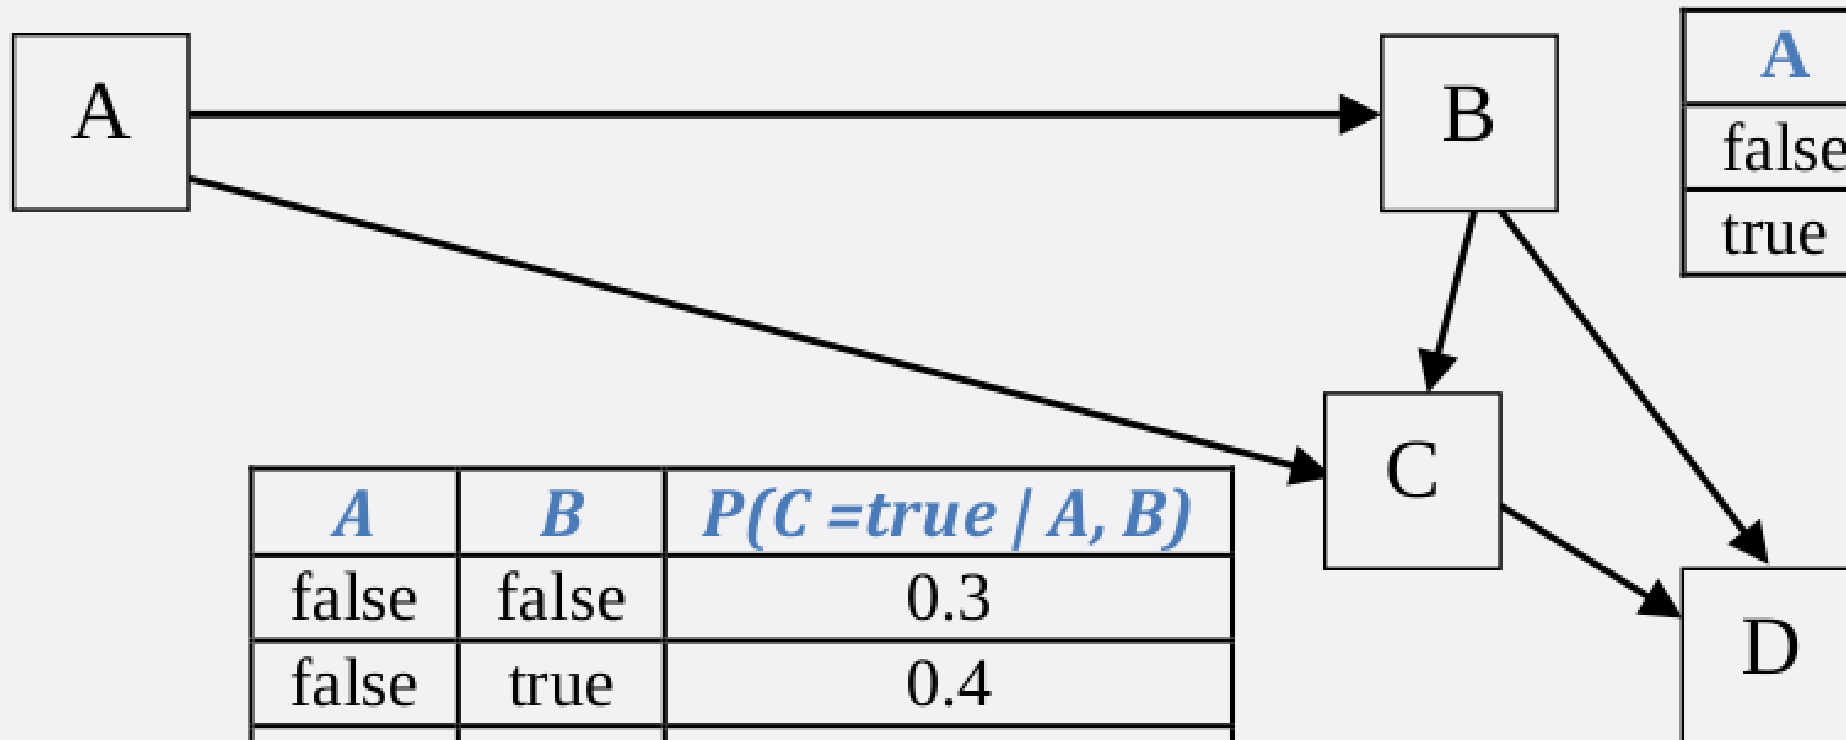

A	P(B =true | A)
false	0.1
true	0.8

A	*B*	*P(C =true | A, B)*
false	false	0.3
false	true	0.4
true	false	0.5
true	true	0.6

B	*C*	*P(D =true | B,C)*
false	false	0.9
false	true	0.3
true	false	0.4
true	true	0.2

2. Compute *P*(B=false and C=false and D=false).

$P(\neg B,\neg C,\neg D) = P(A,\neg B,\neg C,\neg D) + P(\neg A,\neg B,\neg C,\neg D)$
$P(A,\neg B,\neg C,\neg D) = P(A)\ P(\neg B|A)\ P(\neg C|A,\neg B)\ P(\neg D|\neg B,\neg C) = 0.7 * 0.2 * 0.5 * 0.1 = 0.007$
$P(\neg A,\neg B,\neg C,\neg D) = P(\neg A)\ P(\neg B|\neg A)\ P(\neg C|\neg A,\neg B)\ P(\neg D|\neg B,\neg C) = 0.3 * 0.9 * 0.7 * 0.1 = 0.0189$
$P(\neg B,\neg C,\neg D) = 0.007 + 0.0189 = 0.0259$

Exercise 13.3

P(A=true) = 0.7

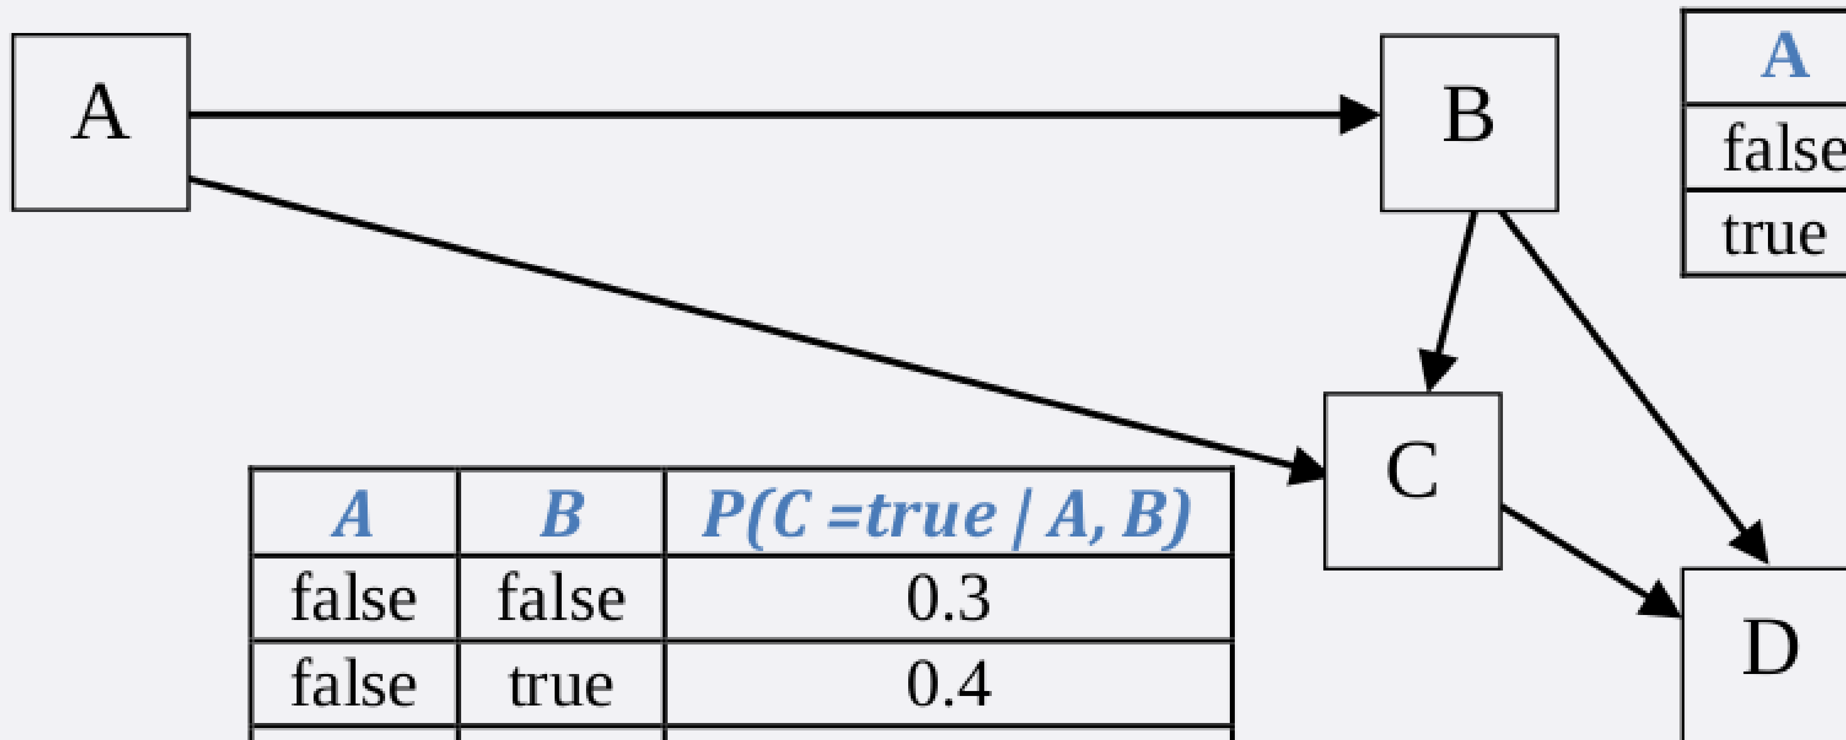

A	P(B =true | A)
false	0.1
true	0.8

A	*B*	*P(C =true | A, B)*
false	false	0.3
false	true	0.4
true	false	0.5
true	true	0.6

B	*C*	*P(D =true | B,C)*
false	false	0.9
false	true	0.3
true	false	0.4
true	true	0.2

3. Compute *P*(A=true | B=true and C=true and D=false).

P(A|B,C,¬D) = P(A,B,C,¬D) / P(B,C,¬D)

P(A,B,C,¬D) = P(A) P(B|A) P(C|A,B) P(¬D|B,C) = 0.7 * 0.8 * 0.6 * 0.8 = 0.2688

P(B,C,¬D) = P(A,B,C,¬D) + P(¬A,B,C,¬D)
P(¬A,B,C,¬D) = P(¬A) P(B|¬A) P(C|¬A,B) P(D|B,C) = 0.3 * 0.1 * 0.4 * 0.8 = 0.0096
P(B,C,¬D) = 0.2688 + 0.0096 = 0.2784

P(A|B,C,¬D) = 0.2688 / 0.2784 = 0.9655

Exercise 13.3

P(A=true) = 0.7

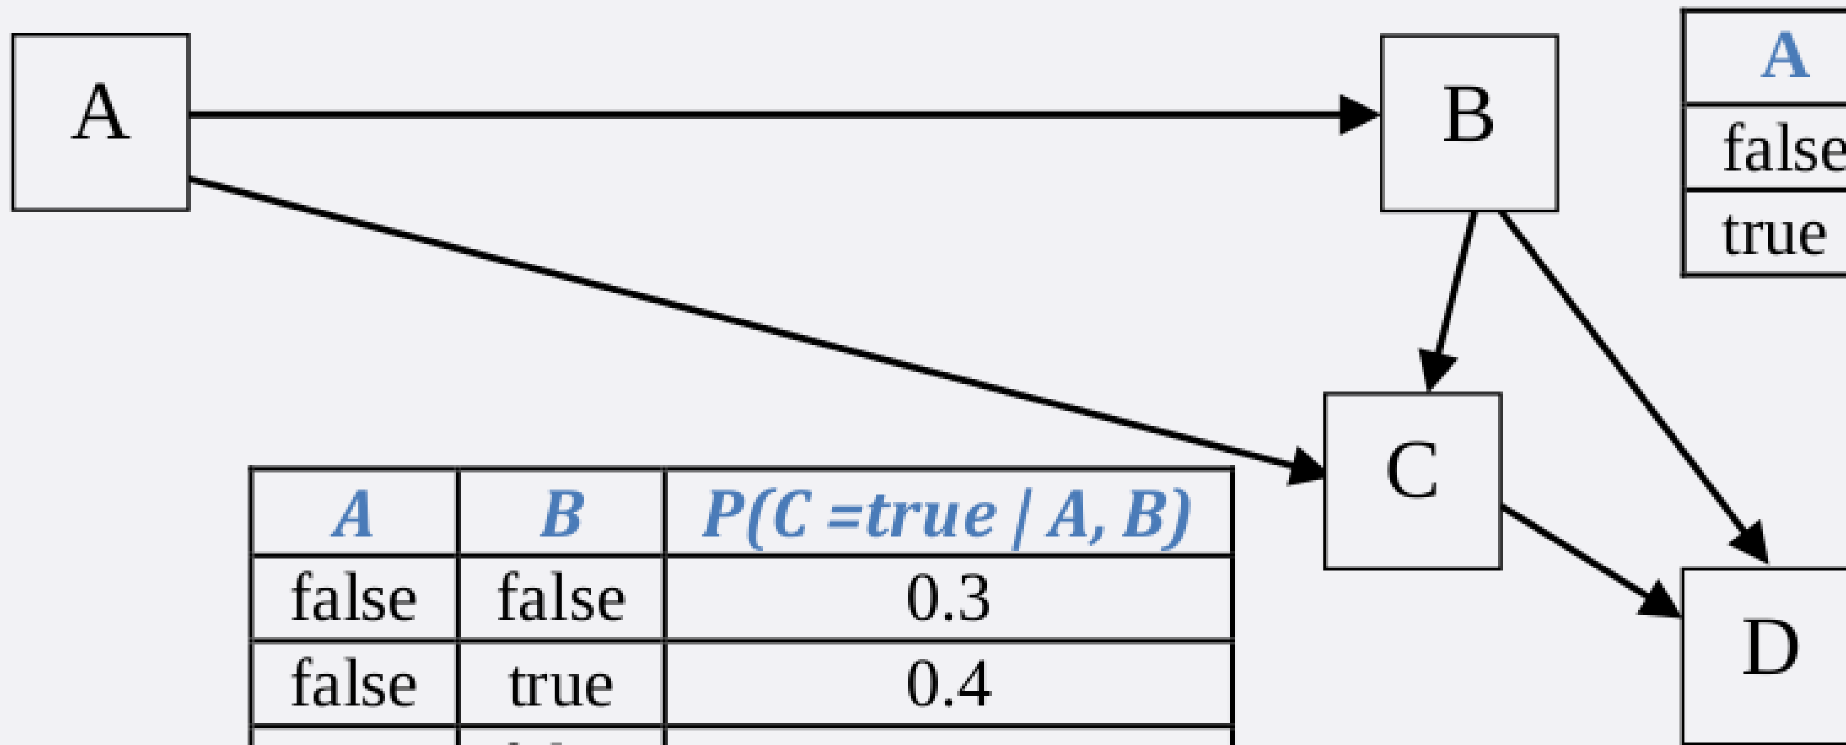

A	P(B =true | A)
false	0.1
true	0.8

A	*B*	*P(C =true | A, B)*
false	false	0.3
false	true	0.4
true	false	0.5
true	true	0.6

B	*C*	*P(D =true | B,C)*
false	false	0.9
false	true	0.3
true	false	0.4
true	true	0.2

3. Compute *P*(A=true | B=true and C=true and D=false).

You could have also observed that due to the Markov Blanket Property, P(A|B,C,¬D) = P(A|B,C) which would have expanded into

P(A|B,C) = P(A,B,C) / P(B,C)
P(A,B,C) = P(A) P(B|A) P(C|A,B) = 0.7 * 0.8 * 0.6 = 0.336
P(B,C) = P(A,B,C) + P(¬A,B,C)
P(¬A,B,C) = 0.3 * 0.1 * 0.4 = 0.012
P(B,C) = 0.336 + 0.012 = 0.348
P(A|B,C) = 0.336 / 0.348 = 0.9655

Exercise 13.3

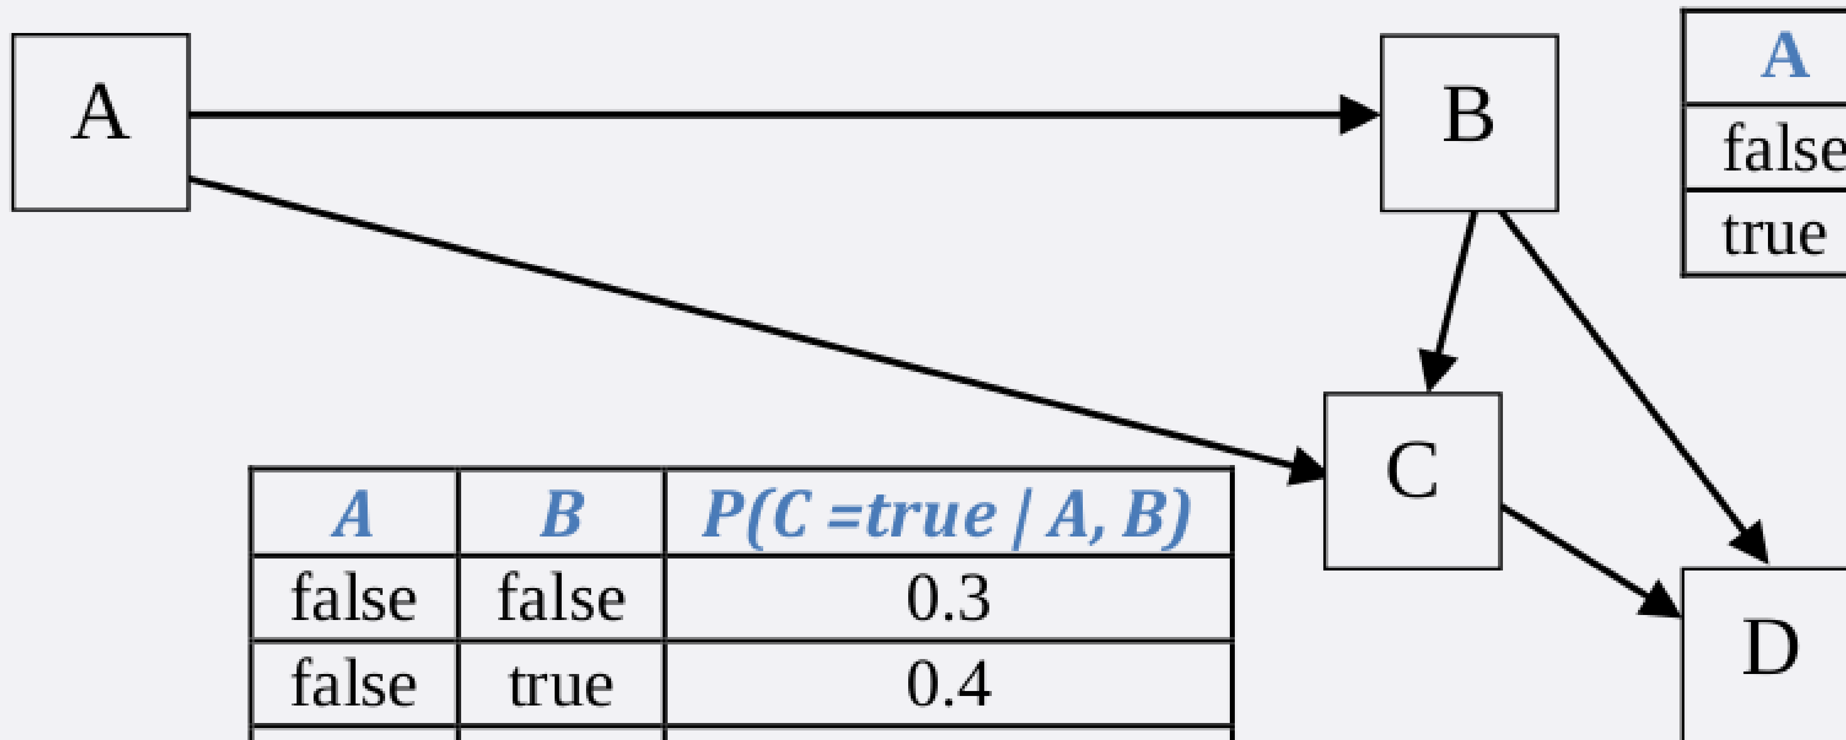

P(A=true) = 0.7

A	P(B =true | A)
false	0.1
true	0.8

A	*B*	*P(C =true | A, B)*
false	false	0.3
false	true	0.4
true	false	0.5
true	true	0.6

B	*C*	*P(D =true | B,C)*
false	false	0.9
false	true	0.3
true	false	0.4
true	true	0.2

4. Compute *P*(D=false | A=true and B=true and C=true).
   (Do this one without employing the Markov Blanket property, i.e., algebraically or numerically confirm that knowing A=true does not change the results from computing *P*(D=false | B=true and C=true).

P(¬D|A,B,C) = P(A,B,C,¬D) / P(A,B,C)
P(A,B,C,¬D) = 0.2688
P(A,B,C) = 0.336
P(¬D|A,B,C) = 0.2688 / 0.336 = 0.8

Exercise 13.3

P(A=true) = 0.7

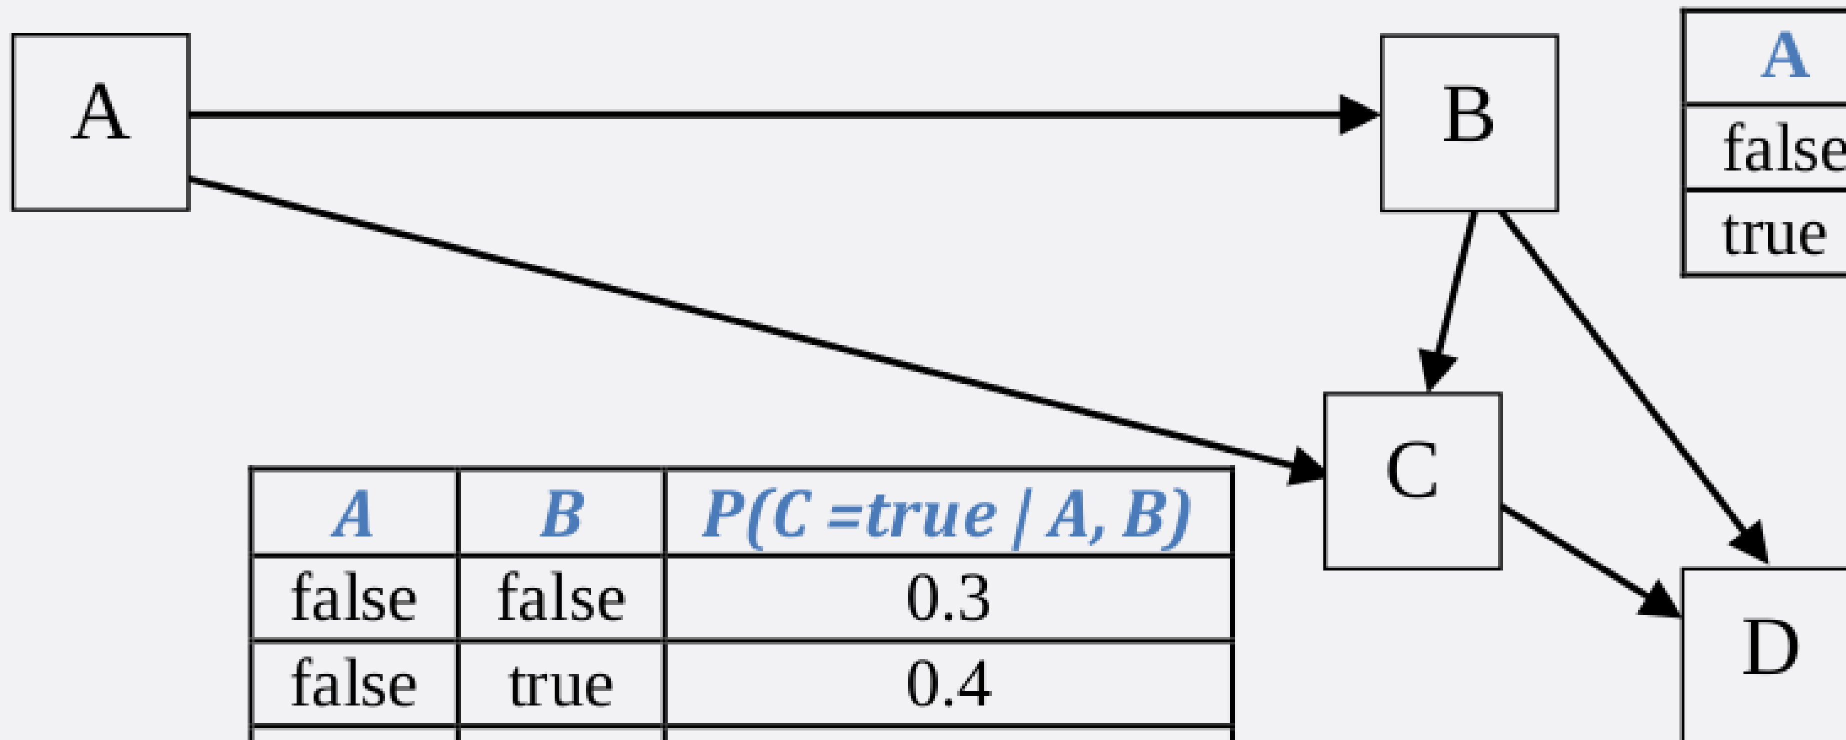

A	P(B =true | A)
false	0.1
true	0.8

A	*B*	*P(C =true | A, B)*
false	false	0.3
false	true	0.4
true	false	0.5
true	true	0.6

B	*C*	*P(D =true | B,C)*
false	false	0.9
false	true	0.3
true	false	0.4
true	true	0.2

4. Compute *P*(D=false | A=true and B=true and C=true).
   (Do this one without employing the Markov Blanket property, i.e., algebraically or numerically confirm that knowing A=true does not change the results from computing *P*(D=false | B=true and C=true).

P(¬D|B,C) = P(B,C,¬D) / P(B,C)
P(B,C,¬D) = 0.2784
P(B,C) = 0.348
P(¬D|B,C) = 0.2784 / 0.348 = 0.8 = P(¬D|A,B,C)

Exercise 13.3

P(A=true) = 0.7

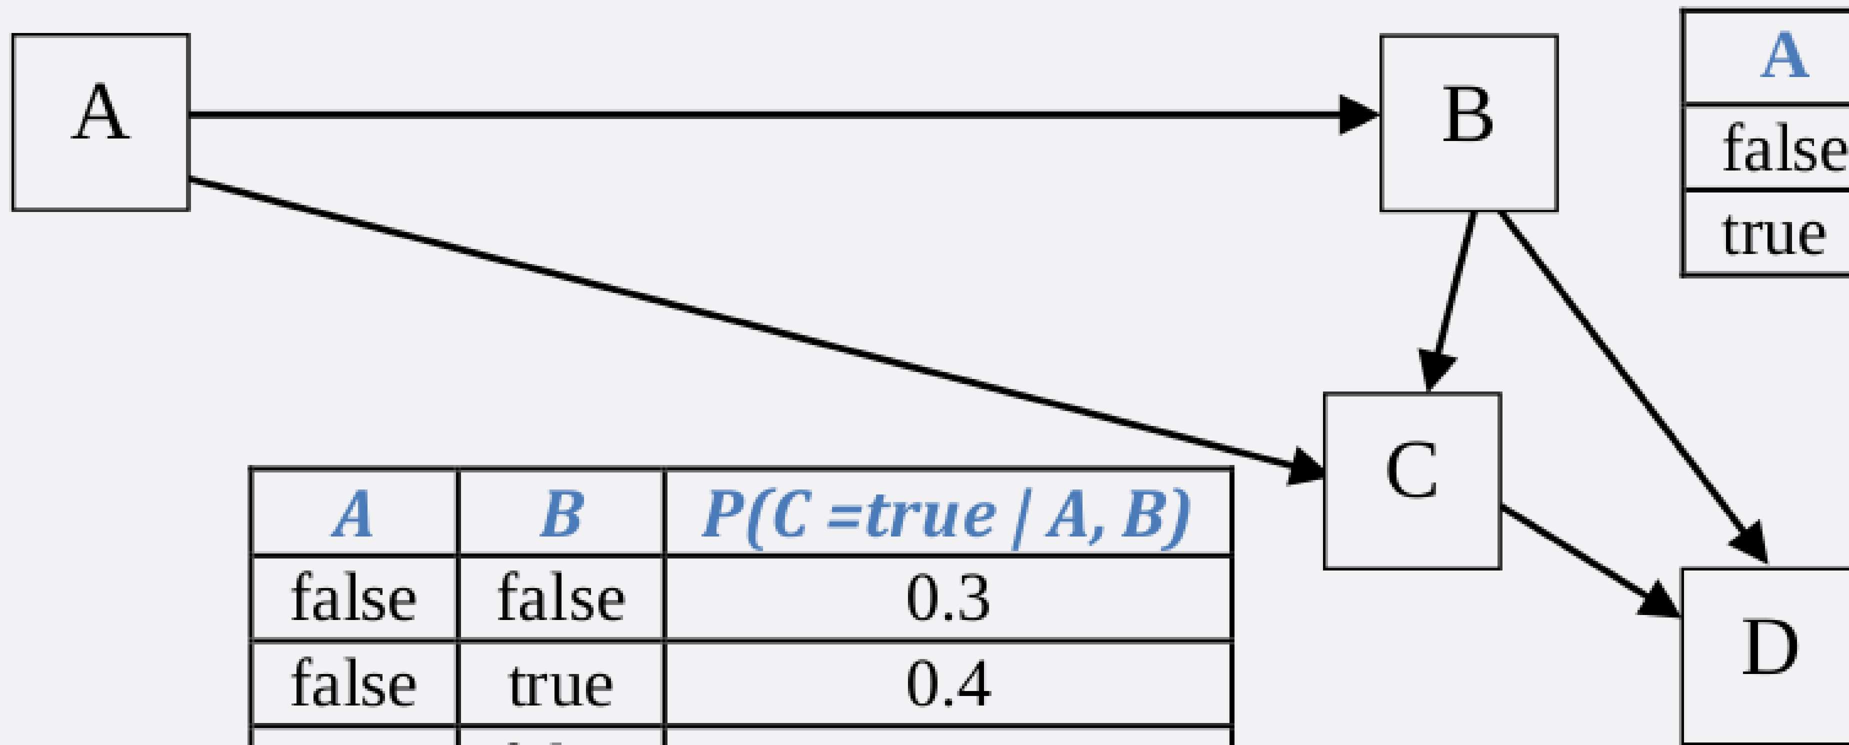

A	P(B =true | A)
false	0.1
true	0.8

A	*B*	*P(C =true | A, B)*
false	false	0.3
false	true	0.4
true	false	0.5
true	true	0.6

B	*C*	*P(D =true | B,C)*
false	false	0.9
false	true	0.3
true	false	0.4
true	true	0.2

5. Compute *P*((A=true and D=true) or (B=true and C=true)).

P((A,D)v(B,C)) = P(A,D) + P(B,C) - P(A,B,C,D)
P(A,D) = P(A,B,C,D) + P(A,¬B,C,D) + P(A,B,¬C,D) + P(A,¬B,¬C,D)
P(A,B,C,D) = 0.0672
P(A,¬B,C,D) = P(A) P(¬B|A) P(C|A,¬B) P(D|¬B,C) = 0.7 * 0.2 * 0.5 * 0.3 = 0.021
P(A,B,¬C,D) = 0.7 * 0.8 * 0.4 * 0.4 = 0.0896
P(A,¬B,¬C,D) = 0.7 * 0.2 * 0.5 * 0.9 = 0.063
P(A,D) = 0.0672 + 0.021 + 0.0896 + 0.063 = 0.2408
P(B,C) = 0.348
P((A,D)v(B,C)) = 0.2408 + 0.348 - 0.0672 = 0.5216

Exercise 13.4

- Your loyal dog Clyde has been howling for the last three hours and you want to decide whether or not to take him to the vet or just to put in ear plugs and go back to sleep. You know that Clyde often howls when there's a full moon, when he's genuinely sick, or occasionally when a particular neighborhood dog starts howling. That neighborhood dog sometimes howls at the full moon and sometimes howls when her owner isn't home, but is not affected by Clyde's howls. If Clyde's really sick he probably won't have eaten very much and should have a bunch of food left in his bowl - but he sometimes just isn't very hungry despite not being sick.

- Question 4.1 - Create a Bayesian network for the scenario described above - use single letter names for each boolean variable (explaining what they mean of course). You can make up the exact numbers in the conditional probability tables (CPTs), but both the CPT values and the causal topology of the network should be reasonable. Briefly explain why you are setting up the topology as you are.

Exercise 13.4

In terms of the CPTs, you just needed to make sure that adding multiple causes increased the probability of a thing - for instance, if the neighbor's dog is howling AND Clyde is sick it should be more probable that Clyde is howling than if only a single one of those causes was true. Additionally, you should have included something around 1/28 for the prior probability of Full Moon (assuming you're on Earth, which is a reasonable assumption).

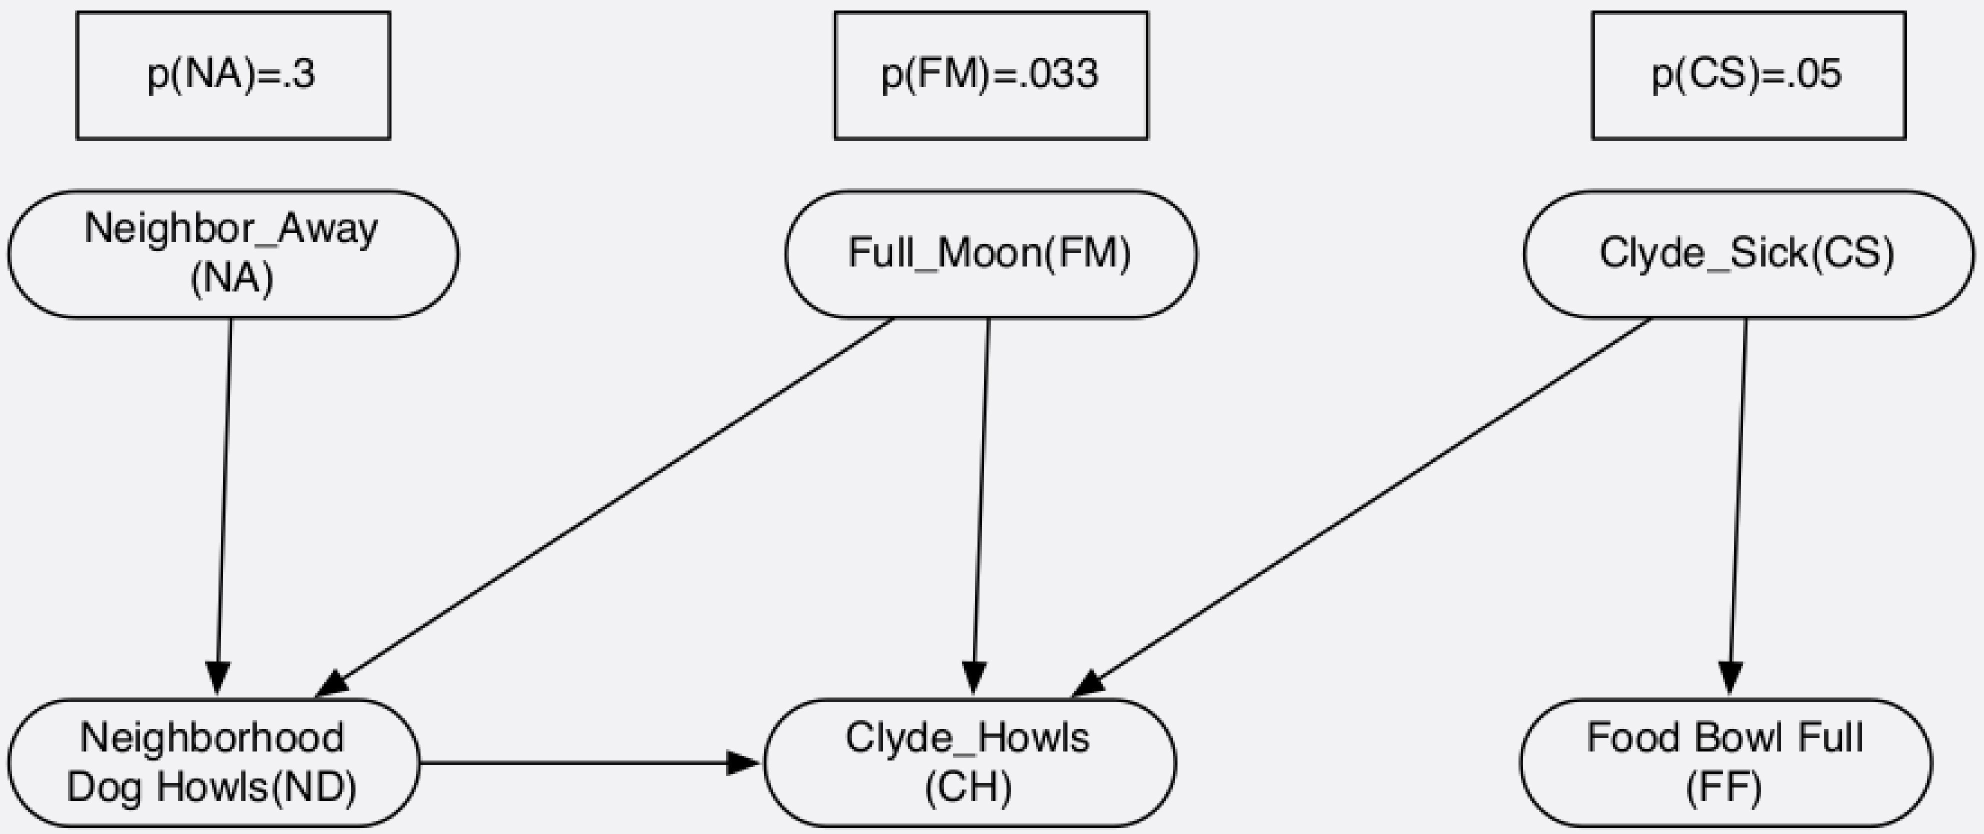

FM	NA	p(ND)
T	T	.8
T	F	.4
F	T	.5
F	F	0

FM	CS	ND	p(CH)
T	T	T	.99
T	T	F	.90
T	F	T	.65
T	F	F	.4
F	T	T	.75
F	T	F	.5
F	F	T	.2
F	F	F	0

CS	p(FF)
T	.6
F	.1

Exercise 13.4

- Your loyal dog Clyde has been howling for the last three hours and you want to decide whether or not to take him to the vet or just to put in ear plugs and go back to sleep. You know that Clyde often howls when there's a full moon, when he's genuinely sick, or occasionally when a particular neighborhood dog starts howling. That neighborhood dog sometimes howls at the full moon and sometimes howls when her owner isn't home, but is not affected by Clyde's howls. If Clyde's really sick he probably won't have eaten very much and should have a bunch of food left in his bowl - but he sometimes just isn't very hungry despite not being sick.

- Question 4.1 - Create a Bayesian network for the scenario described above - use single letter names for each boolean variable (explaining what they mean of course). You can make up the exact numbers in the conditional probability tables (CPTs), but both the CPT values and the causal topology of the network should be reasonable. Briefly explain why you are setting up the topology as you are.
- **Question 4.2 - Identify two conditional independencies in your Bayesian network - express them using the notation A ⊥ B|C. Explain why it makes sense that these variables are conditionally independent given the structure of the problem.**
- **Question 4.3 - For each of the following, write the expression you would need to compute the probabilities using your Bayes Net from 4.1 for inference. You don't need to compute the actual probabilities using your CPT values – just give the expression. Use the notation p(A|B) to indicate the probability that A is true given that B is true - use ¬A to indicate when A is false. Be sure to exploit the structure in your graphs to make the computation easier.**
    1. **What is the probability that Clyde is sick given that there's no full moon, the neighbor's dog is howling, and that his food bowl is full?**
    2. **What is the probability that Clyde is sick given that there's a full moon and you know your neighbor is away, but you have no other information?**

Exercise 13.5

- Would it be rational for an agent to hold the three beliefs?
  - $P(A) = 0.4$
  - $P(B) = 0.3$
  - $P(A \lor B) = 0.5$

- If so, what range of probabilities would be rational for the agent to hold for $A \land B$?
  Make up a table like the one in Figure 13.2, and show how it supports your argument about rationality. Then draw another version of the table where $P(A \lor B) = 0.7$. Explain why it is rational to have this probability, even though the table shows one case that is a loss and three that just break even.

- (Hint: what is Agent 1 committed to about the probability of each of the four cases, especially the case that is a loss?)