Q1. Decision Networks

After years of battles between the ghosts and Pacman, the ghosts challenge Pacman to a winner-take-all showdown, and the game is a coin flip. Pacman has a decision to make: whether to accept the challenge ($accept$) or decline ($decline$). If the coin comes out heads ($+h$) Pacman wins. If the coin comes out tails ($-h$), the ghosts win. No matter what decision Pacman makes, the outcome of the coin is revealed.

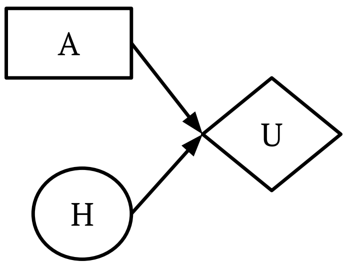

H	$P(H)$
+h	0.5
-h	0.5

H	A	U(H,A)
+h	$accept$	100
-h	$accept$	-100
+h	$decline$	-30
-h	$decline$	50

(a) Maximum Expected Utility

Compute the following quantities:

$EU(accept) = P(+h)U(+h, accept) + P(-h)U(-h, accept) = 0.5 * 100 + 0.5 * -100 = 0$

$EU(decline) = P(+h)U(+h, decline) + P(-h)U(-h, decline) = 0.5 * -30 + 0.5 * 50 = 10$

$MEU(\{\}) = max(0, 10) = 10$

Action that achieves $MEU(\{\}) = decline$

(b) VPI relationships When deciding whether to accept the winner-take-all coin flip, Pacman can consult a few fortune tellers that he knows. There are N fortune tellers, and each one provides a prediction O_n for H.

For each of the questions below, select **all** of the VPI relations that are guaranteed to be true, or select *None of the above*.

(i) In this situation, the fortune tellers give perfect predictions.
Specifically, $P(O_n = +h \mid H = +h) = 1$, $P(O_n = -h \mid H = -h) = 1$, for all n from 1 to N.

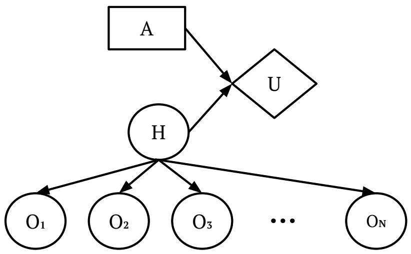

- ☐ $\text{VPI}(O_1,\ O_2) \geq \text{VPI}(O_1) + \text{VPI}(O_2)$
- ■ $\text{VPI}(O_i) = \text{VPI}(O_j)$ where $i \neq j$
- ☐ $\text{VPI}(O_3 \mid O_2,\ O_1) > \text{VPI}(O_2 \mid O_1)$.
- ☐ $\text{VPI}(H) > \text{VPI}(O_1, O_2, \ldots O_N)$
- ☐ None of the above.

(ii) In another situation, the fortune tellers are pretty good, but not perfect.
Specifically, $P(O_n = +h \mid H = +h) = 0.8$, $P(O_n = -h \mid H = -h) = 0.5$, for all n from 1 to N.

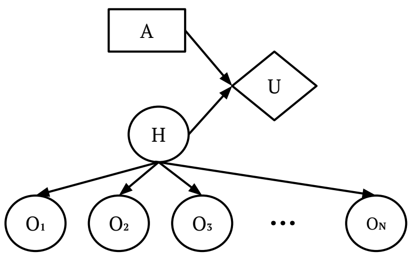

- ☐ $\text{VPI}(O_1,\ O_2) \geq \text{VPI}(O_1) + \text{VPI}(O_2)$
- ■ $\text{VPI}(O_i) = \text{VPI}(O_j)$ where $i \neq j$
- ☐ $\text{VPI}(O_3 \mid O_2,\ O_1) > \text{VPI}(O_2 \mid O_1)$.
- ■ $\text{VPI}(H) > \text{VPI}(O_1, O_2, \ldots O_N)$
- ☐ None of the above.

(iii) In a third situation, each fortune teller's prediction is affected by their mood. If the fortune teller is in a good mood ($+m$), then that fortune teller's prediction is guaranteed to be correct. If the fortune teller is in a bad mood ($-m$), then that teller's prediction is guaranteed to be incorrect. Each fortune teller is happy with probability $P(M_n = +m) = 0.8$.

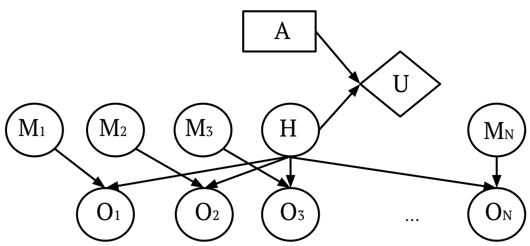

- ☐ $\text{VPI}(M_1) > 0$
- ■ $\forall i\ \ \text{VPI}(M_i|O_i) > 0$
- ☐ $\text{VPI}(M_1,\ M_2,\ \ldots,\ M_N) > \text{VPI}(M_1)$
- ■ $\forall i\ \ \text{VPI}(H) = \text{VPI}(M_i, O_i)$
- ☐ None of the above.

Q2. Probability and Decision Networks

The new Josh Bond Movie (M), Skyrise, is premiering later this week. Skyrise will either be great ($+m$) or horrendous ($-m$); there are no other possible outcomes for its quality. Since you are going to watch the movie no matter what, your primary choice is between going to the theater ($theater$) or renting ($rent$) the movie later. Your utility of enjoyment is only affected by these two variables as shown below:

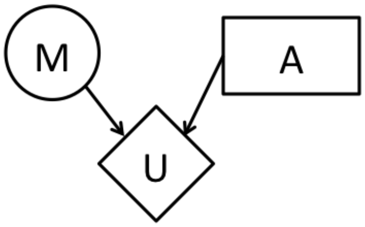

M	P(M)
+m	0.5
-m	0.5

M	A	U(M,A)
+m	*theater*	100
-m	*theater*	10
+m	*rent*	80
-m	*rent*	40

(a) Maximum Expected Utility

Compute the following quantities:

$EU(theater) = P(+m)U(+m, theater) + P(-m)U(-m, theater) = 0.5 * 100 + 0.5 * 10 = 55$

$EU(rent) = P(+m)U(+m, rent) + P(-m)U(-m, rent) = 0.5 * 80 + 0.5 * 40 = 60$

$MEU(\{\}) = 60$

Which action achieves $MEU(\{\}) = rent$

(b) Fish and Chips

Skyrise is being released two weeks earlier in the U.K. than the U.S., which gives you the perfect opportunity to predict the movie's quality. Unfortunately, you don't have access to many sources of information in the U.K., so a little creativity is in order.

You realize that a reasonable assumption to make is that if the movie (M) is great, citizens in the U.K. will celebrate by eating fish and chips (F). Unfortunately the consumption of fish and chips is also affected by a possible food shortage (S), as denoted in the below diagram.

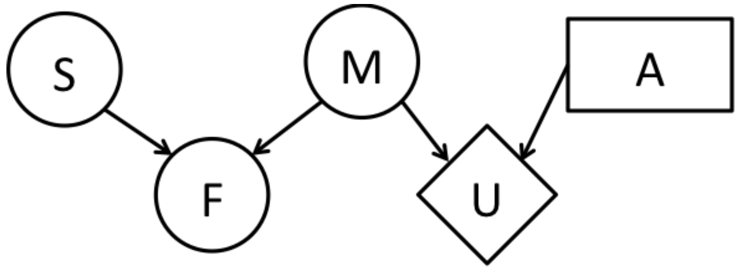

The consumption of fish and chips (F) and the food shortage (S) are both binary variables. The relevant conditional probability tables are listed below:

S	M	F	$P(F\|S,M)$
+s	+m	+f	0.6
+s	+m	-f	0.4
+s	-m	+f	0.0
+s	-m	-f	1.0

S	M	F	$P(F\|S,M)$
-s	+m	+f	1.0
-s	+m	-f	0.0
-s	-m	+f	0.3
-s	-m	-f	0.7

S	$P(S)$
+s	0.2
-s	0.8

You are interested in the value of revealing the food shortage node (S). Answer the following queries:

$EU(theater|+s) =$

The shortage variable is independent of the parents of the utility node when no additional evidence is present; thus, the same values hold:

$EU(theater|+s) = EU(theater) = 55$

$EU(rent|+s) = EU(rent) = 60$

$MEU(\{+s\}) = 60$

Optimal Action Under $\{+s\} = r$ $(Rent)$

$MEU(\{-s\}) = 60$

Optimal Action Under $\{-s\} = r$ $(Rent)$

$VPI(S) = 0$, since the Value of Perfect Information is the expected difference in MEU given the evidence vs. without the evidence and here the evidence is uninformative.

(c) Greasy Waters

You are no longer concerned with the food shortage variable. Instead, you realize that you can determine whether the runoff waters are greasy (G) in the U.K., which is a variable that indicates whether or not fish and chips have been consumed. The prior on M and utility tables are unchanged. Given this different model of the problem:

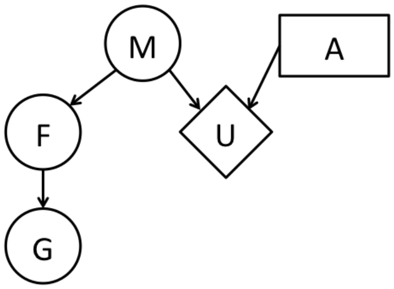

[Decision network]

G	F	$P(G\|F)$
+g	+f	0.8
-g	+f	0.2
+g	-f	0.3
-g	-f	0.7

M	P(M)
+m	0.5
-m	0.5

F	M	$P(F\|M)$
+f	+m	0.92
-f	+m	0.08
+f	-m	0.24
-f	-m	0.76

M	A	U(M,A)
+m	*theater*	100
-m	*theater*	10
+m	*rent*	80
-m	*rent*	40

[Tables that define the model]

F	$P(F)$
+f	0.58
-f	0.42

G	$P(G)$
+g	0.59
-g	0.41

M	G	$P(M\|G)$
+m	+g	0.644
-m	+g	0.356
+m	-g	0.293
-m	-g	0.707

G	M	$P(G\|M)$
+g	+m	0.760
-g	+m	0.240
+g	-m	0.420
-g	-m	0.580

M	F	$P(M\|F)$
+m	+f	0.793
-m	+f	0.207
+m	-f	0.095
-m	-f	0.905

[Tables computed from the first set of tables. Some of them might be convenient to answer the questions below]

Answer the following queries:

$MEU(+g) = max(EU(theater| + g), EU(rent| + g))$

$EU(theater|+g) = P(+m|+g) * U(+m, theater) + P(-m|+g) * U(-m, theater) = (0.644) * 100 + (0.356) * 10 = 67.96$

$EU(rent| + g) = P(+m| + g) * U(+m, rent) + P(-m| + g) * U(-m, rent) = (0.644) * 80 + (0.356) * 40 = 65.76$

$max(EU(theater| + g), EU(rent| + g)) = 67.96$

$MEU(-g) = max(EU(theater| - g), EU(rent| - g))$

$EU(theater|-g) = P(+m|-g) * U(+m, theater) + P(-m|-g) * U(-m, theater) = (0.293) * 100 + (0.707) * 10 = 36.37$

$EU(rent| - g) = P(+m| - g) * U(+m, rent) + P(-m| - g) * U(-m, rent) = (0.293) * 80 + (0.707) * 40 = 51.72$

$max(EU(theater| - g), EU(rent| - g)) = 51.72$

$VPI(G) = P(+g) * MEU(+g) + P(-g) + MEU(-g) - MEU(\{\}) = 0.59 * 67.96 + 0.41 * 51.72 - MEU(\{\}) = 61.3 - 60 = 1.3$

(d) VPI Comparisons

We consider the shortage variable (S) again, resulting in the decision network shown below. The (conditional) probability tables for $P(S)$, $P(M)$, $P(F|S,M)$ and $P(G|F)$ are the ones provided above. The utility function is still the one shown in part (a). Circle all statements that are true, and provide a brief justification (no credit without justification).

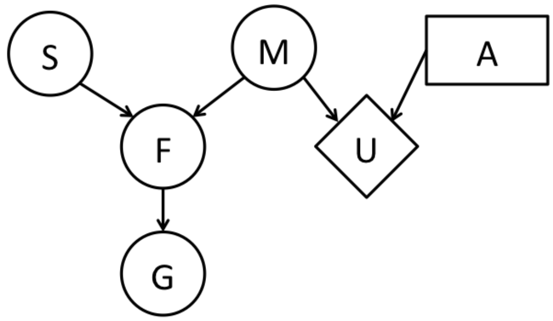

(i) $VPI(S)$:

$VPI(S) < 0$ $\quad$ $VPI(S) = 0$ $\quad$ $VPI(S) > 0$ $\quad$ $VPI(S) = VPI(F)$ $\quad$ $VPI(S) = VPI(G)$

Justify:

With no evidence, VPI(S) is zero because it is conditionally independent of the parents of the utility node.

(ii) $VPI(S|G)$:

$VPI(S|G) < 0$ $\quad$ $VPI(S|G) = 0$ $\quad$ $VPI(S|G) > 0$ $\quad$ $VPI(S|G) = VPI(F)$ $\quad$ $VPI(S|G) = VPI(G)$

Justify:

We accepted several solutions for this question.

Observing G turns $S- > F < -M$ into an active triple, which means S is no longer conditionally independent from the parents of the utility node (M). This introduces the possibility for $VPI(S|G)$ to be strictly positive.

If is possible for $VPI(S|G)$ to be zero due to the presence of conditional independencies that are not exposed by the graph structure (e.g., if F were a coin flip that was independent of S due to a very special choice of CPT entries). It is clear from the CPTs there are no such conditional independencies; consequently, this is not a reason why $VPI(S|G)$ could be 0.

It is possible for $VPI(S|G)$ to be zero if observing S does not change the optimal action for any possible value of S or G (essentially knowing S when G is observed does not allow you to improve your expected performance). Determining this requires substantial computation. We didn't expect you to do so, and accepted justifications that reflect (most of) the above reasoning pattern.

For your interest, we did compute the optimal actions for each situation, and *theater* is optimal for $(+s, +g)$ and for $(-s, +g)$ and *rent* is optimal for $(+s, -g)$ and for $(-s, -g)$ and hence it happens to be the case that $VPI(S|G) = 0$.

(iii) $VPI(G|F)$:

$VPI(G|F) < 0$ $\quad$ $VPI(G|F) = 0$ $\quad$ $VPI(G|F) > 0$ $\quad$ $VPI(G|F) = VPI(F)$ $\quad$ $VPI(G|F) = VPI(G)$

Justify:

G is independent of the parents of the utility node if F is observed, so $VPI(G|F) = 0$

(iv) $VPI(G)$:

$VPI(G) = 0$ $\quad$ $VPI(G) > 0$ $\quad$ $VPI(G) > VPI(F)$ $\quad$ $VPI(G) < VPI(F)$ $\quad$ $VPI(G) = VPI(F)$

Justify:

Since G is a noisy indicator of F, it still has a positive VPI (because VPI of F is positive), but its value is less than the VPI of F. Only if G was an exact representation of F would their VPI's be equal