1 Emily, Car, Stock Market, Sweepstakes, Vacation and Bayes.

Based on the problem statement, we have the Bayesian networks shown in Figure 1 and the corresponding probabilities shown in Table 1.

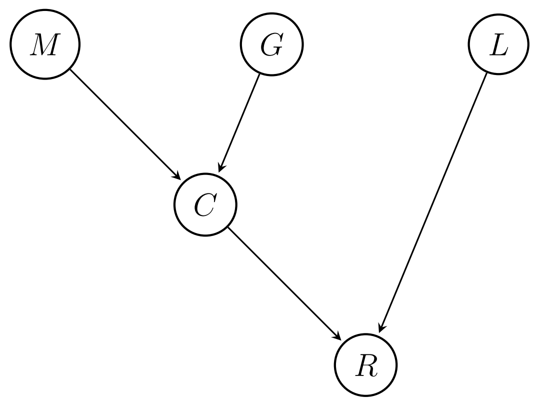

Figure 1: The DAG of the Bayesian networks

Table 1: The known (or elicited) conditional probabilities

Table 2: Market Condition

Market (M)	M^c	M
	0.5	0.5

Table 3: Emily's grade

Grade (G)	A	B	$\leq C$
	0.6	0.3	0.1

Table 4: Car condition

Car (C)	0	1	M	G
	0.5	0.5	M^c	G_A
	0.7	0.3	M^c	G_B
	0.9	0.1	M^c	G_C
	0.2	0.8	M	G_A
	0.5	0.5	M	G_B
	0.8	0.2	M	G_C

Table 5: Probability of Emily went to Redington Shores

Redington Shores (R)	0	1	L	C
	0.8	0.2	L^c	C^c
	0.3	0.7	L^c	C
	0.01	0.99	L	C^c
	0.01	0.99	L	C

We let T^* denote either event T or its complement T^*. Because of Markovian property, the joint probability $P(M^*, G^*, C^*, L^*, R^*)$ can be factorized as

$$P(M^*, G^*, C^*, L^*, R^*) = P(M^*)P(G^*)P(C^*|M^*, G^*)P(L^*)P(R^*|C^*, L^*).$$

We find $P(C|R), P(L|R), P(G_B|R)$ and $P(M^c|R)$ for (a), (b), (c) and (d), respectively. We show the exact calculation for (a) as an example, the detailed computation for (b),(c) and (d) is shown in Appendix A.

(a) We have

$$
\begin{aligned}
P(C,R) &= \sum_{M^*,G^*,L^*} P(M^*,G^*,C,L^*,R) \\
&= P(M)P(G)P(C|M,G)P(L)P(R|C,L) \\
&\quad + P(M)P(G)P(C|M,G)P(L^c)P(R|C,L^c) \\
&\quad + P(M)P(G^c)P(C|M,G^c)P(L)P(R|C,L) \\
&\quad + P(M)P(G^c)P(C|M,G^c)P(L^c)P(R|C,L^c) \\
&\quad + P(M^c)P(G)P(C|M^c,G)P(L)P(R|C,L) \\
&\quad + P(M^c)P(G)P(C|M^c,G)P(L^c)P(R|C,L^c) \\
&\quad + P(M^c)P(G^c)P(C|M^c,G^c)P(L)P(R|C,L) \\
&\quad + P(M^c)P(G^c)P(C|M^c,G^c)P(L^c)P(R|C,L^c) \\
&= 0.3677.
\end{aligned}
$$

By same argument, we have $P(R) = 0.4630$. Hence, we have

$$P(C|R) = \frac{P(C,R)}{P(R)} = \frac{0.3677}{0.4630} = 0.7940.$$

(b)

$$P(L|R) = \frac{P(L,R)}{P(R)} = \frac{0.00099}{0.4630} = 0.0021.$$

(c)

$$P(G_B|R) = \frac{P(G_B,R)}{P(R)} = \frac{0.1202}{0.4630} = 0.2595.$$

(d)

$$P(M^c|R) = 1 - \frac{P(M,R)}{P(R)} = 1 - \frac{0.2627}{0.4630} = 0.4326.$$

Code to the approach of direct simulation using Matlab is shown in Appendix B, and the code to the approach of using OpenBUGS is shown in Appendix C.

2 Trials until Fourth Success.

We first find the posterior distribution based on a general beta prior, namely $\mathcal{B}eta(a,b)$, and n observations.

Let X denote the number of failures until the r-th success in a trail, we know that X follows negative binomial distribution. That is X has pdf as

$$P(X = k) = \binom{k+r-1}{k}(1-p)^k p^r.$$

We substitute $r = 4$ in the above pdf, and obtain $P(X = k) = \binom{k+3}{k}(1-p)^k p^4$. Thus we find the likelihood as

$$L(p|\boldsymbol{x}) = \prod_{i=1}^n p(x_i|p) = \prod_{i=1}^n \binom{x_i+3}{x_i}(1-p)^{x_i} p^4 = \prod_{i=1}^n \binom{x_i+3}{x_i}(1-p)^{\sum_{i=1}^n x_i} p^{4n}.$$

As prior $p \sim \mathcal{B}eta(a, b)$, we have

$$\pi(p) \propto p^{a-1}(1-p)^{b-1}.$$

Then we find the posterior distribution as

$$\begin{aligned}\pi(p|\boldsymbol{x}) &\propto \prod_{i=1}^n \binom{x_i+3}{x_i}(1-p)^{\sum_{i=1}^n x_i} p^{4n} p^{a-1}(1-p)^{b-1} \\ &\propto p^{a+4n-1}(1-p)^{b+\sum_{i=1}^n x_i - 1}.\end{aligned}$$

Thus the posterior follows $\mathcal{B}eta(a + 4n, b + \sum_{i=1}^n x_i)$. As we have $n = 11$ and $\sum_{i=1}^{11} x_i = 30$, the posterior follows $\mathcal{B}eta(a + 44, b + 30)$.

(a) With $a = b = 1$, the posterior is $\mathcal{B}eta(45, 31)$. The Bayes estimator of p is

$$\mathbb{E}_{p|\boldsymbol{x}}[p] = \frac{45}{45+31} = 0.5921.$$

The 95% credible set of p is found as [0.4804, 0.6992] based on the following Matlab command.

```
> betainv(0.025, 45, 31)
[1] 0.4804
> betainv(0.975, 45, 31)
[1] 0.6992
```

The posterior probability of hypothesis $H : p \geq 0.8$ is 0.00001996, computed by the following Matlab command.

```
> 1-betacdf(0.8, 45, 31)
[1] 1.9962e-05
```

(b) With $a = b = 1/2$, the posterior is $\mathcal{B}eta(44.5, 30.5)$. The Bayes estimator of p is

$$\mathbb{E}_{p|\boldsymbol{x}}[p] = \frac{44.5}{44.5 + 30.5} = 0.5933.$$

The 95% credible set of p is found as [0.4809, 0.7011] based on the following matlab command.

```
> betainv(0.025, 44.5, 30.5)
[1] 0.4809
> betainv(0.975, 44.5, 30.5)
[1] 0.7011
```

The posterior probability of hypothesis $H : p \geq 0.8$ is 0.00002479, computed by the following Matlab command.

```
> 1-betacdf(0.8, 44.5, 30.5)
[1] 2.4792e-05
```

(c) With $a = 9, b = 1$, the posterior is $\mathcal{B}eta(53, 31)$. The Bayes estimator of p is

$$\mathbb{E}_{p|\boldsymbol{x}}[p] = \frac{53}{53 + 31} = 0.6310.$$

The 95% credible set of p is found as [0.5257, 0.7303] based on the following matlab command.

```
> betainv(0.025, 53, 31)
[1] 0.5257
> betainv(0.975, 53, 31)
[1] 0.7303
```

The posterior probability of hypothesis $H : p \geq 0.8$ is 0.0001927, computed by the following Matlab command.

```
> 1-betacdf(0.8, 53, 31)
[1] 1.9265e-04
```

3 Penguins.

We first develop Gibbs Sampler that samples from the posterior for μ and τ.

Let Y_i where $i = 1, \dots, n$ denote the measurements of the penguins' height. We know that

$$\begin{aligned} Y_1, \dots, Y_n &\sim \mathcal{N}(\mu, 1/\tau); \\ \mu &\sim \mathcal{N}(\mu_0, 1/\tau_0); \\ \tau &\sim \mathcal{G}a(k, \theta). \end{aligned}$$

where $\mu = 45, \tau_0 = 1/4$, and τ is parameterized by shape parameter k and scale parameter θ. We have $k = 4, \theta = 2$.

The joint distribution is

$$\begin{aligned} \pi(\mu, \tau, \boldsymbol{y}) &= \left\{ \prod_{i=1}^{n} f(y_i|\mu, \tau) \right\} \pi(\mu)\pi(\tau) \\ &\propto \tau^{n/2} \exp\left\{ -\frac{\tau}{2} \sum_{i=1}^{n} (y_i - \mu)^2 \right\} \tau_0^{n/2} \exp\left\{ -\frac{\tau_0}{2}(\mu - \mu_0)^2 \right\} \tau^{k-1} \exp\{-\tau/\theta\} \end{aligned}$$

Thus

$$\begin{aligned} \pi(\mu|\tau, \boldsymbol{y}) &\propto \exp\left\{ -\frac{\tau}{2} \sum_{i=1}^{n} (y_i - \mu)^2 \right\} \exp\left\{ -\frac{\tau_0}{2}(\mu - \mu_0)^2 \right\} \\ &\propto \exp\left\{ -\frac{1}{2}(\tau_0 + n\tau)\left(\mu - \frac{\tau \sum_{i=1}^{n} y_i + \tau_0 \mu_0}{\tau_0 + n\tau}\right)^2 \right\}, \end{aligned}$$

which is a kernel of normal $\mathcal{N}\left(\frac{\tau \sum_{i=1}^{n} y_i + \tau_0 \mu_0}{\tau_0 + n\tau}, \frac{1}{\tau_0 + n\tau}\right)$ distribution. By plugging in the values $n = 14, \sum_{i=1}^{n} y_i = 616, \mu_0 = 45$ and $\tau_0 = 1/4$, we know that $\mu|\tau, \boldsymbol{y} \sim \mathcal{N}\left(\frac{616\tau + 45/4}{1/4 + 14\tau}, \frac{1}{1/4 + 14\tau}\right)$.
Similarly, we have

$$\begin{aligned} \pi(\tau|\mu, \boldsymbol{y}) &\propto \tau^{n/2} \exp\left\{ -\frac{\tau}{2} \sum_{i=1}^{n} (y_i - \mu)^2 \right\} \tau^{k-1} \exp\{-\tau/\theta\} \\ &= \tau^{n/2+k-1} \exp\left\{ -\tau\left(\frac{1}{\theta} + \frac{1}{2} \sum_{i=1}^{n} (y_i - \mu)^2\right) \right\} \end{aligned}$$

which is a kernel of gamma $\mathcal{G}a(k+n/2, 1/(\frac{1}{\theta} + \frac{1}{2}\sum_{i=1}^{n}(y_i - \mu)^2))$ distribution, where the second parameter is a scale parameter. By plugging in the values $n = 14, k = 4$ and $\theta = 2$, we know that $\tau|\mu, \boldsymbol{y} \sim \mathcal{G}a(11, 1/(\frac{1}{2}(1 + \sum_{i=1}^{1} 4(y_i - \mu)^2)))$ where $y_i, i = 1, \dots, 14$ come from the given data. The Matlab code to perform the question (a) and (b) are attached in Appendix D.

(a) The approximated posterior probability of hypothesis $H_0 : \mu < 45$ is 0.9805.

(b) The approximated 95% credible set for τ is $[0.1717, 0.6023]$.

A Matlab code for exact calculation of Emily, Car, Stock Market, Sweepstakes, Vacation and Bayes

```
PR = ...
0.5 * 0.6 * 0.999 * 0.8 * 0.7 +  ...
0.5 * 0.6 * 0.999 * 0.2 * 0.2 +  ...
0.5 * 0.3 * 0.999 * 0.5 * 0.7 +  ...
0.5 * 0.3 * 0.999 * 0.5 * 0.2 +  ...
0.5 * 0.1 * 0.999 * 0.2 * 0.7 +  ...
0.5 * 0.1 * 0.999 * 0.8 * 0.2 +  ...
0.5 * 0.6 * 0.999 * 0.5 * 0.7 +  ...
0.5 * 0.6 * 0.999 * 0.5 * 0.2 +  ...
0.5 * 0.3 * 0.999 * 0.3 * 0.7 +  ...
0.5 * 0.3 * 0.999 * 0.7 * 0.2 +  ...
0.5 * 0.1 * 0.999 * 0.1 * 0.7 +  ...
0.5 * 0.1 * 0.999 * 0.9 * 0.2 +  ...
            0.001 *       0.99

PMR = ...
0.5 * 0.6 * 0.999 * 0.8 * 0.7 +  ...
0.5 * 0.6 * 0.999 * 0.2 * 0.2 +  ...
0.5 * 0.3 * 0.999 * 0.5 * 0.7 +  ...
0.5 * 0.3 * 0.999 * 0.5 * 0.2 +  ...
0.5 * 0.1 * 0.999 * 0.2 * 0.7 +  ...
0.5 * 0.1 * 0.999 * 0.8 * 0.2 +  ...
0.5 *       0.001 * 0.99

PRGB = ...
0.5 * 0.3 * 0.999 * 0.5 * 0.7 +  ...
0.5 * 0.3 * 0.999 * 0.5 * 0.2 +  ...
0.5 * 0.3 * 0.999 * 0.3 * 0.7 +  ...
0.5 * 0.3 * 0.999 * 0.7 * 0.2 +  ...
      0.3 * 0.001 *       0.99

PRC= ...
0.5 * 0.6 * 0.999 * 0.8 * 0.7 +  ...
0.5 * 0.3 * 0.999 * 0.5 * 0.7 +  ...
0.5 * 0.1 * 0.999 * 0.2 * 0.7 +  ...
0.5 * 0.6 * 0.999 * 0.5 * 0.7 +  ...
0.5 * 0.3 * 0.999 * 0.3 * 0.7 +  ...
0.5 * 0.1 * 0.999 * 0.1 * 0.7 +  ...
0.5 * 0.6 * 0.001 * 0.8 * 0.99 +  ...
0.5 * 0.3 * 0.001 * 0.5 * 0.99 +  ...
0.5 * 0.1 * 0.001 * 0.2 * 0.99 +  ...
0.5 * 0.6 * 0.001 * 0.5 * 0.99 +  ...
0.5 * 0.3 * 0.001 * 0.3 * 0.99 +  ...
0.5 * 0.1 * 0.001 * 0.1 * 0.99
```

```
PLR = 0.001 * 0.99

format long
PR  %total probability
PMcgR = 1-PMR/PR
PGBgR = PRGB/PR
PCgR = PRC/PR
PLgR = PLR/PR
format short

%P(R) =0.4630275
%P(Mc|R) = 0.432576898780310
%P(G_B|R) =  0.259546139268186
%P(C|R) =   0.794018173866563
%P(L|R) = 0.002138101948588
```

B Matlab code for simulation of Emily, Car, Stock Market, Sweepstakes, Vacation and Bayes

```
s = RandStream('mt19937ar','Seed',1);
RandStream.setGlobalStream(s);
%
B=100000;
lotteries=[]; cars=[];  markets=[]; gradesB=[]; %save history
redingtonh = 1; %hard evidence
for  i=1:B
lottery=rand≤0.001  ;  % 1 for won, 0 for not won
market = rand ≤ 0.5 ; %1 for bullish, 0 for bearish
grade = mnrnd(1,[0.6, 0.3, 0.1],1) ;
  gradeA=grade(1);
  gradeB=grade(2);
  gradeCorless=grade(3);

if(market)
    if(gradeA)   car=rand≤0.8;
    elseif(gradeB) car=rand<0.5;
    else  car=rand < 0.2;
    end
else
    if(gradeA)   car=rand≤0.5;
    elseif(gradeB) car=rand<0.3;
    else  car= rand < 0.1;
    end
end

if(lottery)
    redington = rand < 0.99;
else
    if(car) redington=rand≤0.7;
    else redington=rand≤0.2;
    end
end


 %%hard evidence filter
         if(redington==redingtonh)
          gradesB=[gradesB gradeB];
          cars=[cars car];
          lotteries=[lotteries lottery];
          markets=[markets market];
          end;
end
  %(a) Got car
```

```
mean(cars)    %0.7931
%(b) Got lottery
mean(lotteries) %0.0022
%(c) Got grade B
mean(gradesB)  %0.2614
%(d) Market bearish
1-mean(markets)  %0.4307
```

C OpenBUGS code for Emily, Car, Stock Market, Sweepstakes, Vacation and Bayes

```
#Model
model {
    market ~ dcat(p.market[])
grade ~ dcat(p.grade[])
gradeA <- equals(grade,1)
gradeB <- equals(grade,2)
    gradeCorless <- equals(grade, 3)
lottery ~ dcat(p.lottery[])
car ~ dcat(p.car[market, grade,])
redington ~ dcat(p.redington[lottery, car, ])
}

#Data
list(redington = 2, p.market=c(0.5, 0.5),
p.grade=c(0.6, 0.3, 0.1),
p.lottery=c(0.999, 0.001),
p.car = structure(.Data = c(0.5, 0.5,      0.7,  0.3,      0.9,  0.1,
                            0.2, 0.8,     0.5,  0.5,      0.8,  0.2),
                            .Dim=c(2,3,2)),
p.redington = structure(.Data = c(0.8, 0.2,       0.3,  0.7,
                                0.01,  0.99,   0.01,  0.99),
                                .Dim=c(2,2,2))  )
```

RESULTS

	mean	sd	MC_error	val2.5pc	median	val97.5pc	start	sample
car	1.793	0.4051	6.443E-4	1.0	2.0	2.0	1001	1000000
gradeB	0.2592	0.4382	5.886E-4	0.0	0.0	1.0	1001	1000000
lottery	1.002	0.04572	8.137E-5	1.0	1.0	1.0	1001	1000000
market	1.567	0.4956	6.533E-4	1.0	2.0	2.0	1001	1000000

(a) P(car|redington)=0.793
(b) P(won lottery|redington) = 0.002
(c) P(grade B|Redington) = 0.2592
(d) P(market bearish\redington) = 1-0.576 = 0.424

D Matlab code for Penguins

```
%FALL 2019 --  MIDTERM Online Course ISyE6420 (Penguins)
%full conditional distributions available
%
%  y_i ¬ N(mu, 1/tau), i=1,...,n
%  mu  ¬ N(mu0, 1/tau0); mu0=45, tau0=1/4
%  tau ¬ Ga(a,1/b);  shape=4, rate=1/2
%------------------------------------------
clear all;
close all;
clc;
%-----------------figure defaults
lw = 2;
set(0, 'DefaultAxesFontSize', 17);
fs = 14;
msize = 5;
%--------------------------------
 n=14; % sample size
 randn('state', 10);
 x=[41 44 43 47 43 46 45 42 45 45 43 45 47 40];
 suma = sum(x);
%------------------------------------------
%
nn = 10000+1000;
mus = [];
taus = [];
mu = 40; tau =8; % start with the chain the parameters as prior means
mu0=45; tau0=1/4;
h=waitbar(0,'Simulation in progress');
for i = 1 : nn


  new_mu   = normrnd( (tau * suma+tau0*mu0)/(tau0+n*tau),   ...
     sqrt(1/(tau0+n*tau)) );
  par      = 1/2 + 1/2 * sum ( (x - mu).^2 );
  new_tau =gamrnd(4 + n/2, 1/par);
  mus = [mus new_mu];
  taus = [taus new_tau];
  tau=new_tau;
  mu=new_mu;
 if i/50==fix(i/50)            % Shows wait bar
    waitbar(i/nn)
 end
end
 close(h)
%
burnin = 1000;
```

```
figure(1)
subplot(2,1,1)
plot((nn-burnin:nn), mus(nn-burnin:nn))
xlabel('Mu')
subplot(2,1,2)
plot((nn-burnin:nn), taus(nn-burnin:nn))
xlabel('Tau')

figure(2)
subplot(2,1,1)
hist(mus(burnin:nn), 70)
title('Mu');
subplot(2,1,2)
hist(taus(burnin:nn), 70)
title('Tau');

figure(3)
plot( mus(burnin:nn), taus(burnin:nn), '.')
xlabel('Mu');
ylabel('Tau');
title('Scatter plot of new mu and new tau');
mean(mus(burnin:nn))  %44.0477
mean(taus(burnin:nn))  %0.3557

 mean(1./taus(burnin:nn))   %3.1095
 %
 %Posterior mean using gibbs sampler
 %After burnin 500 records
 [mean(mus(burnin:nn)) std(mus(burnin:nn))  prctile(mus(burnin:nn),2.5) ...
    median(mus(burnin:nn)) prctile(mus(burnin:nn),97.5)]
    %  44.0493    0.4604   43.1308   44.0504   44.9601

 %Posterior precision (1/sig2) using gibbs sampler
  %After burnin 500 records
 [mean(taus(burnin:nn)) std(taus(burnin:nn))  ...
    prctile(taus(burnin:nn),2.5) median(taus(burnin:nn)) ...
    prctile(taus(burnin:nn),97.5)]
     %0.3554    0.1095    0.1735    0.3452    0.5995

%(a)
sum(mus(burnin:nn)  < 45)/(nn-burnin)
%0.9805
%(b)
[prctile(taus, 2.5)  prctile(taus, 97.5)]
% 0.1717    0.6023
```