Final. Closed book and calculators not allowed. Answers may include e^2, $\sqrt{\ }$, etc. but simplfy when possible.

Your Name: ____________________

Problem 1.
Recall that a Gaussian prior is conjugate to the mean of a Gaussian distribution.
Given:
1. a random variable X is distributed normally given its mean, i.e. $X|\mu \sim N(\mu, 1)$ 2. our prior belief regarding μ is a standard normal: $\mu \sim \mathcal{N}(0, 1)$
3. we have one data point $x_1 = 10$.

Question: what is the posterior distribution of μ after observing x_1?

1a. Informally justify your answer (可以用中文)

1b. (Challenging?) Mathematically prove your answer.

Your Name: ____

Problem 2.
Imagine rolling a (not necessarily fair) 4-sided die, numbered {1,2,3,4}.
Given:
1. prior: Your prior belief on the probability of each side is Dirichlet($\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}$).
2. data: You roll the die twice, with getting a '1' and '3'.

Question:
2a. What is the posterior distribution over {1,2,3,4} after observing the data?

2b. What is the probability that the next die roll yields a 3?

Your Name: ____________________

Problem 3.

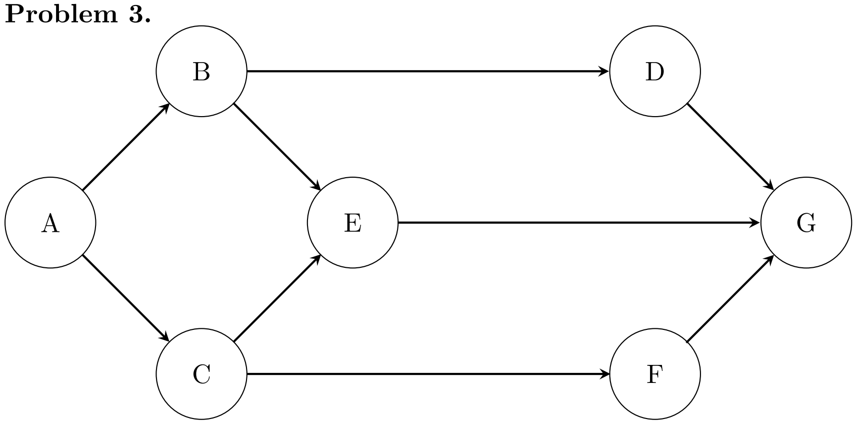

The above graph is a Bayesian network (aka Belief Network, or probabilistic graphical model). Consider the $\binom{7}{3} = 35$ possible triples of nodes (A,B,C); (A,B,D); ... ; (E,F,G).

Question:
List the triples (X,Y,Z) for which X and Y are conditionally independent given Z.
Where $\mathrm{X}, \mathrm{Y}, \mathrm{Z} \in \{\mathrm{A}, \dots, \mathrm{G}\}, \mathrm{X} \neq \mathrm{Y}, \mathrm{X} \neq \mathrm{Z}, \mathrm{Y} \neq \mathrm{Z}$.

Your Name: ______________________________

Problem 4.

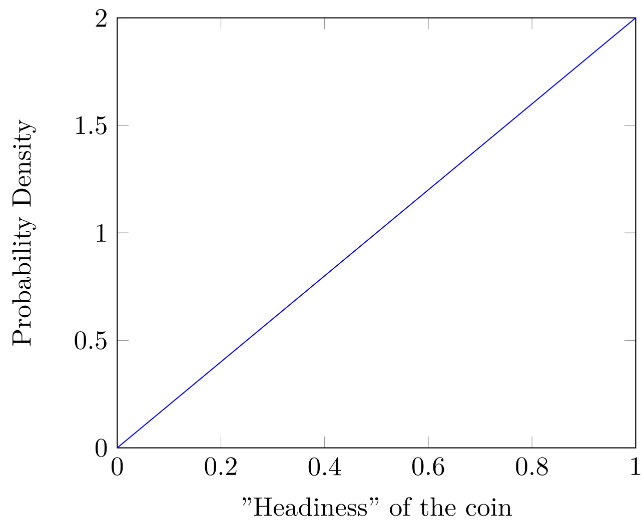

This is a coin flipping problem.
Recall a beta distribution is defined as:

$$\text{BetaDist}(a, b) \stackrel{\text{def}}{=} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}\, p^{a-1}\,(1-p)^{b-1}$$

Given:
1. the data is a single coin toss, yielding "heads".
2. a beta distribution $\text{BetaDist}(a, b)$ was used as a prior.
3. the posterior distribution is as plotted above.

Question:
What were the parameters (a, b) of the beta distribution prior?

Your Name: ______________________________

Problem 5.
Dataset:

Class	F_1	F_2	F_3
A	good	good	okay
A	bad	bad	good
A	bad	okay	okay
A	okay	okay	good
A	bad	okay	good
B	good	okay	okay
B	okay	okay	bad
B	okay	good	bad
B	good	bad	bad

Question:
Specify a Naïve Bayes classifer based on the above dataset.
Your classifier should provide enough information to compute the numerical value of $\mathrm{P}[class = \mathrm{A}|\mathrm{F}_1, \mathrm{F}_2, \mathrm{F}_3]$ for all 9 combinations of $(\mathrm{F}_1, \mathrm{F}_2, \mathrm{F}_3) \in \{\mathrm{good}, \mathrm{okay}, \mathrm{bad}\}$.
Explicitly state all priors used.