Bayesian Modelling B sheet 1

1. Two rational mathematics students, A and B, play the following game. An unknown amount of money, £w, is placed in an envelope, £$2w$ is placed in another envelope, and both are sealed. The envelopes are then shuffled, and one given to A and one to B. Before they each open their envelopes, A is given the option of switching envelopes with B, if B agrees. In either case, each student keeps the money in the envelope he opens.

   A argues that "There is an unknown sum x in my envelope. The other envelope is equally likely to contain $2x$ or $x/2$. On average, therefore, I will get $(2x + x/2)/2 = 1.25x$ if I switch, which is more than x. Therefore I should switch." B says the expected gain from switching is zero.

   Who is right? Use probability to explain.

2. On the "Monty Hall" TV game show (in the US in the 70's), a contestant is shown 3 doors. Behind one is a valuable prize (say, a car), behind the other two something worthless (traditionally, a goat). The host of the show knows where the car is. The contestant, who does not know, chooses a door to be opened. Before that door is opened, the host opens one of the *other* doors, and exposes a goat. The contestant is given an option to reconsider her choice (among the two still closed). She thinks that this option gives her an even chance of getting the car, instead of the 2 to 1 odds against that she started with. Should she switch her choice? Or is it true that the host is giving her no information, since whatever door she chooses initially, there remains one that he can open to show a goat?

   There is a lot about this problem on the Internet; some of what you read there is correct. One fun page is `http://math.ucsd.edu/~crypto/Monty/monty.html`

3. Suppose that for three random variables x, y, z, the joint distribution factorises as a product of a function of x and z, and a function of y and z, i.e. $p(x, y, z) = f(x, z)g(y, z)$ (where f and g are arbitrary functions, not assumed to be p.d.f.'s). It is claimed in the notes that this implies $x \perp y \mid z$. Prove this claim for discrete random variables.

4. Let $x_0, x_1, \ldots, x_n$ be part of a realisation of a Markov chain. Show that the joint distribution of $(x_0, x_1, \ldots, x_n)$ factorises as a product of functions of adjacent pairs of chain values (x_{i-1}, x_i). Conversely, if the joint distribution of $(x_0, x_1, \ldots, x_n)$ factorises in this way, show that $(x_{i+1}, x_{i+2}, \ldots, x_n) \perp (x_0, x_1, \ldots, x_{i-1}) \mid x_i$ for all i, using the result of the previous question.

5. Consider a variant of the '10+1' coin tossing problem from Section 1, where instead of a discrete choice between 2 biased coins, the parameter θ is supposed to be drawn from a Beta(α, β) prior. Write down the joint distribution $p(\theta, x, y)$. Integrate out θ to find $p(x, y)$. Indicate how you would use this to find the conditional expectation $E(y|x)$.

   Note that you get the answer much more easily by first finding the posterior $p(\theta|x)$, and then noting that $E(y|x) = P\{y = 1|x\} = E(\theta|x)$ (which we already know, or can easily find).

6. Suppose that, given λ, x and y are independent Poisson(λ). Show that the distribution of x given that $x + y = n$ is Binomial$(n, 0.5)$. Hence show that $\lambda \perp x \mid x + y$. More generally, if $x_1, x_2, \ldots, x_m$ are i.i.d. Poisson(λ), given λ, then $\lambda \perp (x_1, x_2, \ldots, x_m) \mid \sum_{i=1}^m x_i$, that is, if we observe the value of $\sum_{i=1}^m x_i$, looking at the individual values of the $\{x_i\}$ tells us nothing more about λ: we say $\sum_{i=1}^m x_i$ is *sufficient* for λ.