Simulating Language, week 3 pre-reading questions, part 3

These last few questions are about putting likelihoods and priors together to do inference. Remember, we are going to be modeling language as a process of inferring a grammar based on data and a prior probability distribution over possible language types - so this dice example might seem a bit weird and irrelevant, but the idea is it's underlyingly exactly the same kind of inference that's involved.

1. Now let's put all the bits together. Imagine you have a large bag containing 50 dice. 49 of those dice are 'fair' dice, equally likely to roll 1, 2, 3, 4, 5 or 6. However, the 50th dice is loaded, and *always* rolls a 6. You friend is going to dip their hand into the bag, pull out a dice, roll it a few times, and you have to use Bayes' Rule to decide if it is a normal dice or the loaded dice.

   a. Let's build up the various component parts. First, what is the prior probability that your friend will pull a normal, 'fair' dice out of the bag? In the notation used in the reading, we would represent this probability as something like p(fair-dice).

   b. What is the prior probability of them pulling out the loaded dice [i.e. p(loaded-dice)]?

   c. If they pull out a fair dice, what is the probability of getting a 6 when they roll that dice? In the notation used in the reading, we could write this down as p(6 | fair- dice), i.e. the probability of rolling a 6 given that we are rolling a fair dice.

   d. If they pull out a fair dice, what is the probability of getting two 6s in a row when they roll that dice, i.e. p(6,6 | fair-dice)?

   e. If they pull out the loaded dice, what is the probability of getting a 6 when they roll that dice, i.e. p(6 | loaded-dice)?

   f. If they pull out the loaded dice, what is the likelihood of getting two 6s in a row when they roll that dice, i.e. p(6,6 | loaded-dice)?

   g. Now, you friend reaches into the bag, pulls out a dice, and rolls it - it's a 6! Then they roll it again - another 6! Using Bayes' Rule: is it more probable that they are rolling a fair dice or the loaded dice?

   Hint: Bayes Rule states that the posterior probability of a hypothesis given some data is proportional to the prior probability of that hypothesis times the likelihood of the data given that hypothesis. To get actual posterior probabilities, you have to divide by the probability of the data, but all that does is normalise everything, so you don't need to do it to answer this question - you can figure out the answer by just considering the priors and likelihoods that you worked out above. The hypotheses are that the dice is fair or it is loaded; the data is the dice rolls.

   h. Now your friend rolls the dice a third time. Another 6! Based on this new data, is it more probable that they are rolling a fair dice or the loaded dice?