FINAL EXAM – submit by September 10th

Instructor: David Freeborn

Total: 120 points

Question 1: (6 points) The following statements are either true or false. Circle the correct answer.

(a)	Every sound argument is also valid.	True	False
(b)	Some inductive arguments are valid.	True	False
(c)	Every tautology has probability 1.	True	False
(d)	Two mutually exclusive propositions *can* be true at the same time.	True	False
(e)	Every chance set-up with *independent* trials has biased outcomes.	True	False
(f)	Two propositions are independent if the truth of one does not affect the probability of the other.	True	False

Question 2: (6 points) Two events A and B are *independent*. Suppose that $\Pr(A) = 0.1$ and $\Pr(B) = 0.3$. What is the probability that:

(i) neither A nor B occurs?

(ii) at least one of these events occurs?

(iii) exactly one (i.e. not both) of these events occurs?

(iv) Are the two events mutually exclusive? How do you know?

Question 3: (10 points) Suppose we have 4 apples (a, b, c, d) and 5 buckets $(1, 2, 3, 4, 5)$. We place each apple in a random bucket; the placement of each apple is independent of the others. Let's call B_{ij} the proposition that apples i and j are in the same bucket. For example, B_{ab} is the proposition that apples a and b are in the same bucket. That could happen in 5 ways: they could both be in bucket 1, they could both be in bucket 2, they could both be in bucket 3, they could both be in bucket 4, or they could both be in bucket 5. For each question, explain your reasoning as well as giving an answer.

(a) (2 points) What is the probability of each proposition, B_{ij}?

(b) (2 points) Think carefully. Is B_{ab} independent of B_{cd}?

(c) (2 points) Is B_{ab} independent of B_{bc}?

(d) (2 points) What is the probability of $B_{bc} \wedge B_{cd}$?

(e) (2 points) Is B_{ab} independent of $B_{bc} \wedge B_{cd}$?

Question 4: (12 points) Consider the following chance setup. You have two identical six-sided dice with non-standard marking. Each is unbiased with respect to its six sides. **For each of them, three of the sides have the number 1 written on them; two of the sides have the number 2 written on them; and the final side has the number 3**. Answer each of the following questions (a) - (d).

(a) (2 points) If you roll ***one*** of the dice, what is the probability of yielding each of the possible outcomes?

(b) (2 points) If you roll ***one*** of the dice, what is the probability of yielding an odd number?

(c) (4 points) If you roll ***both*** dice, what is the probability that their sum will be an even number?

(d) (4 points) If you roll ***both*** dice, what is the probability of yielding each of the possible outcomes for their sum?

Question 5: (14 points) I toss an unbiased coin **twice** in a row. If I get heads both times, I draw from urn A. Otherwise I draw from urn B.

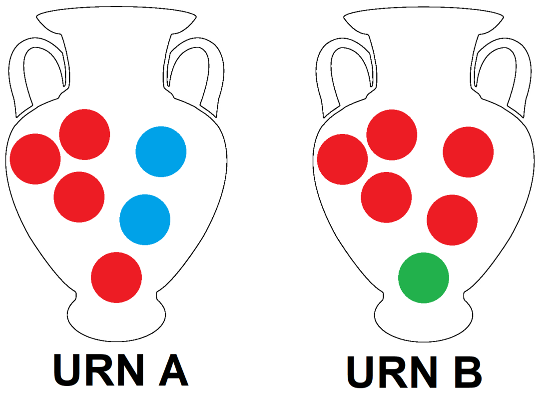

(a) (2 points) What is the probability of drawing a red ball?

(b) (1 point) Suppose that I told you that I drew a green ball. What is the probability that the coin landed heads twice?

(c) (2 points) Suppose that I told you that I drew a green ball. What is the probaiblity that the coin landed heads on the first roll?

(d) (2 points) Suppose that I told you that I drew a green ball. What is the probaiblity that the coin landed heads at least once?

(d) (7 points) Suppose that I told you that I drew a red ball. What is the probaiblity that the coin landed heads at least once?

Question 6: (15 points) A certain disease affects about 1 out of every 500 people. Fortunately, there is a test that is very sensitive to the disease: if you have the disease, the test is positive 99% of the time. If you do not have the disease, then there is a 98% chance that you will get a negative result on the test. Use Bayes' theorem and the total probability theorem to calculate the following.

a) (3 points) For a person selected a random from the whole population, what is the probability of their blood producing a positive result on the test?

b) (3 points) What is the probability that someone has the disease, ***given*** that she got a positive result on the test?

c) (4 points) Assume that successive tests are statistically independent. Suppose a patient took the test twice. What is the probability that she has the disease, ***given*** that she got positive results both times?

d) (4 points) Suppose the patient took the a third time. The third time, the patient got a negative result. What is the probability that she has the disease, ***given*** these test results?

e) (1 point) In the assumption that successive tests are independent reasonable?

Question 7: (16 points) Ben and Jerry's is hiring a team of nutritionists to test whether a new diet, the *Nothing But Ice Cream* (NBIC) diet, to see if it improves health outcomes. The nutritionists will claim the diet works if they can attain a p-value of < 0.05.

a) (1 point) What is the null hypothesis in this experiment?

b) (3 points) What does it mean to attain a p-value of < 0.05? You may want to include a diagram in your explanation.

c) (4 points) The scientists carry out 20 different statistically independent trials. Let's suppose that, unfortunately, in the real world, *Nothing But Ice Cream* diet does not work. What is the probability that *none* of the trials show a positive result?

d) (4 points) In fact, exactly one trial shows a positive result. What was the probability that exactly one trial would show a positive result?

.

e) (2 points) The well-known journal, *Dedicated Experimental Literature on Ice Cream*

Investigative Operations Undertaking Science (DELICIOUS), only publishes the positive outcome. What is the problem with this?

f) (2 points) Suggest some better scientific practices than the ones undertaken by the nutritionists in this question.

Question 8: (12 points) Breit Bouchard thinks that he has understood Professor Freeborn's lectures on probability theory. "Bayes' rule tells us how to update our priors based on evidence", he says, "but we are free to choose whatever prior probabilities that we like. Freedom is the American way!"

a) (3 points) Breit assigns the following prior probabilities. The probability that he can recycle a randomly selected shampoo bottle is 0.7. And he believes that half of all shampoo bottles are made from high density polyethene (HDPE). And if a shampoo bottle is recyclable, then there is a 90% chance that it is made from HDPE.

Breit learns that his shampoo bottle is in fact made from HDPE. Put these numbers into Bayes' theorem to find his updated probability that his shampoo bottle can be recycled. What's the problem with the answer?

b) (1 point) Look back at the Kolmogorov axioms. Which axiom does this violate?

c) (3 points) "You lied to me!" Breit yells at Professor Freeborn. "I put my beliefs into Bayes' theorem, and I got out a number that's not even a probability!"

Professor Freeborn shakes his head in disappointment. "For someone so interested in waste disposal, I'm surprised you don't know the dictum, 'Garbage in, garbage out'. Your initial probabilities were inconsistent. Bayes' theorem will only produce a sensible answer if you put in consistent probabilities."

Breit looks dumbfounded and heartbroken. Can you explain why the initial probabilities were inconsistent? There are lots of ways to explain this, and drawing a diagram (optional) can help with some explanations.

d) (5 points) Let's help Breit choose some consistent initial beliefs. He is sure that that half of all shampoo bottles are made from high density polyethene (HDPE). And he is sure that, if a shampoo bottle is recyclable, then there is a 90% chance that it is made from HDPE. Given these two beliefs, what is the highest possible starting probability he can assign to any particular bottle being recyclable (if he doesn't know anything more about what it is made from)?

Question 9: (15 points) The detective Sherlock Holmes is searching for the thief who has stolen Brianne's homework. There are two suspects, Professor Moriarty (P) and Colonel Moran (C). Sherlock Holmes' priors are that there is a 50% chance each of them could have committed the crime. And there are two forms of evidence available: DNA and witnesses.

a) (4 points) If Colonel Moran committed the crime, Sherlock Holmes thinks that the witnesses would correctly identify him 90% of the time. But Professor Moriarty is sneaky and able to fool witnesses. If Professor Moriarty did the crime, Sherlock Holmes thinks that witnesses would correctly identify him only 50% of the time. Given Sherlock Holmes thinks it is equally likely that either Moriarty or Moran, how likely is it that a witness will give the correct result?

b) (2 points) DNA evidence is right about $\frac{2}{3}$ of the time, regardless of who committed the crime. In this case, which is the better form of evidence, DNA or witnesses? Why?

c) (2 points) Do you think we can treat witnesses and DNA as independent sources of evidence?

d) (4 points) In fact, both the witness and the DNA evidence suggest that Professor Moriarty committed the crime. Supposing that they are independent sources of evidence, how probable is it that Professor Moriarty is the homework thief?

e) (1 point) Suppose that instead of DNA evidence, Sherlock Holmes had a second witness. Do you think that a second witness be a better or worse form of evidence than DNA evidence? Why or why not? (No right/wrong answer – you can argue it either way.)

f) (2 points) Holmes' assistant, Watson, has been spending time at the opium den. He makes the following claim: "In the limit that we had not just two witnesses, but infinitely many witnesses to the crime, then we would know with certainty who committed the crime. Simply interview all the witnesses, and accept their majority verdict. That verdict would be certain to be correct." Is Watson right?

Question 10: (14 points) Prior to the discovery of Australia, Europeans generally believed that all swans were colored white. Indeed, the phrase "as rare as a black swan" was a popular saying (quoting the Roman poet Juvenal) to refer to an impossible event. This is because every swan that European people had ever seen was white. At the end of the 18th century, European explorers encountered black swans for the first time, in the continent of Australia.

a) (4 points) Suppose we consider two hypotheses, H_1 "all swans are white", and H_2 "half of swans are white, half of swans are black". Let us start with a 50% probability in each hypothesis. Let's suppose we pick some swan at random, from the population of all swans. It turns out this swan is white. How should we update our probabilities in each hypothesis?

b) (1 point) Let's suppose that we keep drawing white swans. Without calculation, what will happen to the our degree of belief in H_1 and H_2 if we keep conditionalizing using Bayes' rule?

c) (1 point) Is the assumption that we draw a swan at random from the population of all swans, a good model of what Europeans were doing before the discovery of Australia?

d) (3 points) Let's suppose the 1000th swan I draw from the population is neither white nor black. It is pink with Polka dots, with a tartan-patterned beak and legs. Uh oh. I didn't even consider that possibility. In principle, can I accommodate this result with Bayes' rule? What "rule of thumb" does this relate to?

e) (2 points) Suppose that instead we start with the two hypotheses, H_a all swans are white, H_e "all swans in Europe are white but swans elsewhere are black". Each has prior probability $\frac{1}{2}$. Now suppose that we see a European white swan. How should we update our conditional beliefs in each hypothesis?

f) (1 point) Why do you think Europeans did not consider the hypothesis H_e when they observed white swans in Europe?

g) (2 points) Professor Froopey proposes, "Instead of the categories "black" and "white", Europeans should have considered two different categories:

- *blite*: An object is blite if it is black and we observe it outside of Australia, or it is white and we observe it in Australia."
- *whack*: An object is whack if it is white and we observe it outside of Australia, or it is black and we observe it in Australia."

"These would have been better categories to make inductive generalizations," says Professor Froopey. "Every white swan we observe in Europe, or black swan that we observe in Australia is a whack swan. So the simplest and clearest inductive generalization to make is that "all swans are whack"."

Do you agree with Professor Froopey?

End of Exam! Well done! Remember to fill out the evaluations, and enjoy the rest of the summer!