CS 70 Final 9 August 2012

NAME (1 pt): ______________________________

SID (1 pt): ______________________________

TA (1 pt): ______________________________

Name of Neighbor to your left (1 pt): ______________________________

Name of Neighbor to your right (1 pt): ______________________________

Instructions: This is a closed book, closed calculator, closed computer, closed network, open brain exam, but you are permitted a 3 page, double-sided set of notes, large enough to read without a magnifying glass.

You get one point each for filling in the 5 lines at the top of this page.

Write all your answers on this exam. If you need scratch paper, ask for it, write your name on each sheet, and attach it when you turn it in (we have a stapler).

1	
2	
3	
4	
5	
6	
7	
8	
Total	

Question 1 (20 points) True or False.

For each of the following propositions, circle either T if it is always true, F if it is always false. You do not have to justify your answer.

We let $\mathbb{N} = \{0, 1, 2, ...\}$ denote the non-negative integers, $\mathbb{Q}$ denote the set of rational numbers. A and B denote events in the sample space S.

T F For all propositions P, $[P(0) \wedge P(1) \wedge (\forall k \in \mathbb{N} P(k) \implies P(k+2))] \implies \forall n \in \mathbb{N} P(n)$
Answer: *True*

T F Here P and Q denote propositions, each of which could be true or false:
$[P \implies (Q \wedge \neg Q)] \implies \neg P$
Answer: *True*

T F Given a list of x values for a polynomial and the delta functions generated from Lagrange Interpolation, they can be used to pass through an arbitrary combination of y values.
Answer: *True*

T F There are a total of p^{d+1} polynomials of degree d in $GF(p)$
Answer: *True*

T F For a uniformly distributed random variable defined on $[a, b]$, it's sufficient to know the expectation and variance of the random variable to uniquely determine the intervals a and b.
Answer: *True*
For uniform distribution on x, $E(x) = \frac{a+b}{2}, V(x) = \frac{(b-a)^2}{2}$. We have two equations and two unknowns.

T F There are $\binom{n+k}{k}$ different ways to throw k identical balls into n distinguishable bins. Each bin can have multiple balls.
Answer: *False*
There are $\binom{n+k-1}{k}$ different ways.

T F Markov's inequality, Chebyshev and the Central Limit Theorem all are useful for upper bounding hard to compute probabilities.
Answer: *False*

T F The multiplicative inverse of a modulo M exists $\iff$ M is prime.
Answer: *False*
a and M coprime

T F The set of degree n polynomials over $\mathbb{Q}$ is countable
Answer: *True*

T F The set of numbers of the form $(\sqrt{x} + \sqrt{y})^2$ where $x, y \in \mathbb{Q}$ is not countable
Answer: *False*

Question 2 (25 points)

2.1 (5 points). Here A and B denote propositions, each of which could be true or false. Prove or disprove the following proposition and clearly state whether the proposition is true or false.

$$((\neg A \implies B) \land (\neg A \implies \neg B)) \implies A$$

Answer:

$$\begin{aligned}
((\neg A \implies B) \land (\neg A \implies \neg B)) \implies A \\
&= ((A \lor B) \land (A \lor \neg B)) \implies A \\
&= \neg((A \lor B) \land (A \lor \neg B)) \lor A \\
&= \neg(A \lor B) \lor \neg(A \lor \neg B) \lor A \\
&= (\neg A \land \neg B) \lor (\neg A \land B) \lor A \\
&= \neg A \lor A \\
&= True
\end{aligned}$$

2.2 (5 points). Prove that among any set of n+1 integers one can find 2 numbers so that their difference is divisible by n.

Answer: Prove by pigeonhole. There are at least 2 numbers $i \neq j$ where $x_i \equiv x_j \pmod{n}$. Therefore $x_i - x_j \equiv 0 \pmod{n}$ and $x_i - x_j$ is divisible by n.

2.3 (5 points). It's possible to run a version of the traditional propose & reject algorithm where the men and women switch roles (women propose, men say "maybe" or "yes", etc.). We'll refer to this as the non-traditional propose & reject algorithm. Prove the following statement:

In a stable marriage instance, there is exactly one stable pairing **if and only if** the pairing produced by the traditional propose & reject algorithm is the same as the pairing produced by the non-traditional propose & reject algorithm.

Answer: To prove $P \iff Q$ we need to prove $P \implies Q$ and $Q \implies P$

Let's start by proving the following: if there is exactly one stable pairing then the pairing produced by the traditional marriage is the same as the pairing produced by the non-traditional marriage algorithm

If there is exactly one stable pairing we'll call that pairing T. Because TPRA and NTPRA must output a stable pairing, they will both output T.

Now for the hard part: if the pairing produced by the traditional marriage is the same as the pairing produced by the non-traditional marriage algorithm then there is exactly one stable pairing.

Assume they both give the same pairing (T) but there are multiple stable pairings. We know that in the pairing T, every pair (m_i, w_i) is optimal for both partners. Now, let's look at another "stable pairing", T' where some couple is difference. In this case, without loss of generality, we have the couples (m_i, w_j) and (m_k, w_i). Based on the definition of optimal, we know that m_i prefers w_i to his wife, w_j and w_i prefers m_i to her husband m_k. Therefore T' is unstable which contradicts the premise that it is a stable pairing different from T.

2.4 (10 points). Prove the following identity by induction:

$$\forall n \geq 0. \sum_{i=1}^{n+1} i * 2^i = n * 2^{n+2} + 2$$

a) State and prove the base case for the induction.

Answer:
$n = 0$

$$1 * 2^1 = 0 * 2^2 + 2 = 2$$

b) State the induction hypothesis.

Answer:
Assume for some $k \in \mathbb{N}$, $\sum_{i=1}^{k+1} i * 2^i = k * 2^{k+2} + 2$.

c) Complete the proof by stating and proving the induction step.

Answer:
Prove $\sum_{i=1}^{k+2} i * 2^i = (k+1) * 2^{k+3} + 2$

$$\begin{aligned}
\sum_{i=1}^{k+2} i * 2^i &= \sum_{i=1}^{k+1} i * 2^i + (k+2) * 2^{k+2} \\
&= k * 2^{k+2} + 2 + (k+2) * 2^{k+2} \\
&= 2^{k+2}(k + k + 2) + 2 \\
&= 2^{k+2}(2k + 2) + 2 \\
&= 2^{k+2}2(k+1) + 2 \\
&= (k+1) * 2^{k+3} + 2
\end{aligned}$$

Question 3 (15 points)

3.1 (1 points). If it exists, find the inverse of: 20 mod 79 and express your answer as the smallest valid non-negative integer. **Answer:** 4

3.2 (2 points). If it exists, find the inverse of: 5 mod 23 and express your answer as the smallest valid non-negative integer. **Answer:** 14

3.3 (1 points). What is $2^{2^{2006}}$ mod 3? Express your answer as the smallest valid non-negative integer. **Answer:**

$$2^{2^{2006}} \equiv -1^{2^{2006}} \equiv 1 \pmod 3$$

3.4 (1 points). Calculate $2^{125} \mod 127$. Express your answer as the smallest valid non-negative integer.

Answer:

$$\begin{aligned} 2^{125} &\equiv 2^{-1} 2^{126} \\ &\equiv 2^{-1} \\ &\equiv 64 \pmod{127} \end{aligned}$$

3.5 (5 points). Suppose that p is prime. Let X be an integer uniformly distributed on $\{1, \ldots, p^n - 1\}$. What is the probability that X doesn't have an inverse modulo p^n?

Answer: Multiples of p are not coprime with p^n. So, all numbers of the form ap, $a \in \{1, 2, \ldots, p^{n-1} - 1\}$. So, $p^{n-1} - 1$ numbers don't work. Thus, the desired probability is $\frac{p^{n-1}-1}{p^n - 1}$.

3.6 (5 points). Suppose that p is prime. Let X be a geometric random variable with probability 1/2 of a successful trial. What is the probability that X doesn't have an inverse modulo p^n? (hint: $\sum_{i=1}^{\infty} a^i = \frac{1}{1-a} - 1$ when $0 < a < 1$)

Answer: Multiples of p are not coprime with p^n. So, all numbers of the form ap, $a \in \{1, 2, \ldots\}$. So

$$\begin{aligned} &\Pr(X \text{ doesn't have an inverse}) \\ &= \sum_{i=1}^{\infty} \Pr(X \text{ doesn't have an inverse} \cap \{X = i\}) \\ &= \sum_{i=1}^{\infty} I\{i \text{ doesn't have an inverse}\} \frac{1}{2}^i \\ &= \sum_{i=1}^{\infty} I\{i = kp \text{ for some } k\} \frac{1}{2}^i \\ &= \frac{1}{2}^p + \frac{1}{2}^{2p} + \ldots \\ &= \sum_{j=1}^{\infty} [\frac{1}{2}^p]^j \\ &= \frac{1}{1 - 0.5^p} - 1 \end{aligned}$$

Question 4 (15 points) Some people just don't care

Let's suppose we have a stable marriage instance with n pairs of men and women. For whatever reason, everyone has decided that the other sex is pretty much the same to them, so rather than create a preference list by ranking, they will generate a list by ordering the other sex uniformly at random. Assume each person's list is independently generated.

4.1 Warm up (5 points). You do not need to show your work for these (they are worth 1 point each).

(a) How many different preference lists can a person generate?
Answer: $n!$

(b) What is the probability that M_1 will have W_1 at the top of his list?
Answer: $\frac{(n-1)!}{n!} = \frac{1}{n}$

(c) What is the probability that *TMA* will terminate in exactly 1 day?
Answer: For *TMA* to terminate in exactly 1 day, each man must have a different woman at the top of his list. The probability that M_2 has a different first choice from M_1 is $\frac{n-1}{n}$, the probability that M_3 has a difference first choice from M_1 and M_2 is $\frac{n-2}{n}$, which means our probability that all men have different first choices is

$$\frac{n!}{n^n}$$

(d) What is the probability that *TMA* will terminate in exactly 1 day and produce a female-optimal pairing?
Answer: For each set of men's preference lists that list a different woman first, the probability that W_1 will have the man who put her first at the top of her list is $\frac{1}{n}$. Since each woman independently generates her list, the probability that every woman will have the "correct" man listed first is $(\frac{1}{n})^n$. This means our final probability is

$$\frac{n!}{n^n} \cdot \frac{1}{n^n} = \frac{n!}{n^{2n}}$$

Note: Technically, this is not the only way that *TMA* will terminate in exactly 1 day and produce a female-optimal pairing. However, this is the solution that we accepted (and the one most people gave–no one got the general case, because it is rather difficult to compute).

(e) What is the probability that every man has the same preference list?
Answer: $(\frac{1}{(n!)})^{n-1}$

4.2 (5 points). What is the probability that M_1 will prefer W_1 over W_2? Explain your answer.
Answer 1: For each preference list where W_1 is ahead of W_2, we can switch their positions, and we will have a preference list where W_2 is ahead of W_1. Therefore, the probability that we will get a preference list where W_1 is ahead of W_2 is simply $\frac{1}{2}$.

Answer 2: Consider the case where M_1 places W_1 at the i^{th} position in his list. The number of ways to place W_2 after W_1 is

$$\begin{aligned}
|\text{ways to place } W_2 \text{ after the } i^{th} \text{ position}| &= \frac{(n-2)!}{(n-2-(i-1))!} \cdot (n-i)! && (1)\\
&= \frac{(n-2)!}{(n-i-1)!} \cdot (n-i)! && (2)\\
&= (n-i) \cdot (n-2)! && (3)
\end{aligned}$$

There are two ways to count this:
(1) says we first pick the women to go in slots 1 through $i-1$ (this is $n-2$ permute $i-i$), then the ways to place the rest of the women (this is $(n-i)!$).
(3) says to first place W_2 (there are $n-i$ choices for this), then place the rest of the women (there are $n-2$ women to be placed).

This means that the total number of ways for W_1 to be in front of W_2 is

$$\begin{aligned}
|\text{ways to place } W_2 \text{ after } W_1| &= \sum_{i=1}^{n-1}[(n-i)\cdot(n-2)!]\\
&= (n-2)! \cdot \sum_{i=1}^{n-1}(n-i) = \sum_{i=1}^{n-1} i\\
&= (n-2)! \cdot \frac{(n-1)n}{2}\\
&= \frac{n!}{2}
\end{aligned}$$

Finally, our probability is

$$\begin{aligned}
\Pr(M_1 \text{ likes } W_1 \text{ better } W_2) &= \frac{\frac{n!}{2}}{n!}\\
&= \frac{1}{2}
\end{aligned}$$

4.3 (5 points) After randomly generating his preference list, M_1 suddenly realizes that he does have a preference for the women after all! However, he's still not sure what that preference is, and so he comes up with the following way to figure it out:

1. Order the pairs of positions on his list (note that this is a pair of **positions**, *not* women). Denote each position in a pair as r_1 and r_2.
   For each pair:
2. If he likes the woman currently listed at r_2 more than r_1, switch the women's places. Otherwise, leave them in their current positions.
3. Repeat until all pairs have been checked.

What is the expected number of switches he will have to make before he recovers his preference list? Assume he never changes his mind about two women.
[Hint: Each swap decreases the number of inverted pairs by 1.]

Answer: Let S be a random variable denoting the number of swaps he will have to make, and let S_{ij} be a random indicator variable denoting whether or not r_i and r_j need to be switched (ie he prefers the woman listed in position r_j to the woman listed in position r_i). We see that $S = \sum S_{ij}$, where $i < j$. Furthermore, from 8.2, we know that $\Pr(S_{ij}) = \frac{1}{2}$. Lastly, the total number of indicator variables is simply $\binom{n}{2}$, since he will only need one swap to correctly place each woman. Therefore,

$$\begin{aligned}
\mathbb{E}(S) &= \mathbb{E}(\sum S_{ij}) \\
&= \sum \mathbb{E}(S_{ij}) \\
&= \sum \Pr(S_{ij} = 1) \\
&= \sum \frac{1}{2} \\
&= \binom{n}{2} \times \frac{1}{2}
\end{aligned}$$

Question 5 (15 points) Caverns and Drakes

After much anticipation, the wonderful people at Hail® have released their newest RPG (role-playing game), *Caverns and Drakes*. Not wanting to fall behind your friends, who camped out to buy it, you quickly read up on the rules and decide to create a powerful fighter.

5.1 (5 points). Character creation works as follows: roll 4 fair six-sided dice, ignore the lowest roll (in the event of a tie, pick one to ignore), and total the remaining results (for example, if the dice resulted in 6, 5, 5, 3, the total would be $6 + 5 + 5 = 16$). Do this twice to determine your Attack and Defense ratings.

(a) **2 points.** Identify the probability space ω.
Answer: We could define an outcome ω as (a, b, c, d, t), where a denotes the result of rolling the first die, b denotes the result of rolling the second die, c denotes the result of rolling the third die, d denotes the result of rolling the fourth die, and t denotes the result of the total (add the results, ignoring the lowest). However, this is not a uniform distribution (for example, the outcome $(1, 1, 1, 1, 18)$ would have probability 0).

But because t is completely determined by a, b, c, and d, there are a total of 6^4 possible outcomes with non-zero probability, and each has probability $\frac{1}{6^4} = \frac{1}{1296}$ of occurring.
Similarly, we can instead define an outcome as (a, b, c, d), and we would have a uniform distribution.

Note: Students who defined the probability space as t invariably got stuck when assigning probabilities to each outcome. This is because, in general, the probabilities are very difficult to calculate. We were hoping that parts (b) and (c) would give a hint that defining the space as t was not correct, since otherwise you would be repeating work.

(b) **2 points.** What is the probability of getting an Attack rating of 18?
Answer: To roll an 18, you need a result of 6 on 3 of the dice, and the last die can take on any value. The number of ways to achieve this is $\binom{4}{3} \cdot \binom{5}{1} + 1 = 21$ (first pick which 3 dice will be a 6, then pick the value of the remaining die. We separate out the case where the last die is a 6, since we would otherwise be overcounting).
The probability is therefore $\frac{21}{1296} \approx 1.6\%$.

(c) **1 point.** What is the probability of getting an Attack rating of 3?
Answer: There is only one way to have the total be 3, and that is when the result of all 4 dice is 1.
Thus, the probability is $\frac{1}{1296} \approx 0.08\%$.

5.2 (5 points). Let's suppose that after you roll, your Attack and Defense are both 10. It's time to go out and kill things!

Combat works as follows: each round, roll a 20-sided die, add your Attack rating to the result (we will call this the *Attack Value*), and compare it to your target's Defense + 10 (the extra 10 on Defense is for balance issues–otherwise, everyone would be hitting everything!). If you exceed that number, you hit and deal damage. Otherwise, you miss (and deal no damage).
Each character (and monster) also has a Hit Point total, which damage is subtracted from. If your Hit Point total goes below 0, you die.

You must show your work–answers with no justification will receive *no* credit.

Answer: Many students were confused by the schematic (for example, some tried to add the Attack Value to damage, or tried to use the damage as the Attack Value). We accepted solutions that made these assumptions (eg giving 17 as an answer to 1(b)).

(a) **1 point.** A Troll has a Defense of 10. What is the probability that you will hit it with an attack?
Answer: The number to beat is 20, meaning that any result of 11 or higher on the die will succeed. Therefore, the probability is simply $\frac{10}{20} = 0.5$.

(b) **1 point.** Suppose you have a greatsword, whose damage is calculated as follows: roll two 6-sided dice and add the results. Given that you hit the Troll, how much damage can you expect to deal?
Answer: This is just the expected value of the sum of two 6-sided dice. Because of linearity of expectation, it is the same as twice the expected value of one die. Formally, if D_1 is the result of the first die roll, and D_2 is the result of the second die roll, $\mathbb{E}(D_1 + D_2) = \mathbb{E}(D_1) + \mathbb{E}(D_2) = 3.5 + 3.5 = 7$.

(c) **1 point.** How much damage should you expect to deal to the Troll each round?
Answer: Let H be the random variable denoting whether or not the sword hits. On a success, it will deal 7 damage, otherwise it will deal 0 damage. From (e), we know

$$Pr[H = \text{hit}] = \frac{10}{20}$$
$$Pr[H = \text{miss}] = \frac{10}{20}$$

Thus, $\mathbb{E}(H) = \frac{10}{20} \cdot 7 + \frac{10}{20} \cdot 0 = 3.5$.

(d) **2 points.** If you encounter a Troll, should you fight it?
(A Troll has an Attack rating of 14 and 40 Hit Points. You have 50 Hit Points. The damage for a Troll's attack, its claw, is calculated as follows: roll a 6-sided die and add 6.5 to the result.)
Answer: Since your Defense is 10, the Troll will hit you with probability $\frac{14}{20}$ (it only needs a result of 7 or higher to beat 20). Furthermore, its expected damage if it hits you is $3.5 + 6 = 10$. This means that, each round, you can expect it to deal $0.7 \cdot 10 = 7$ damage to you. You should expect it to take you more than 10 rounds to down the Troll, but only about 7 rounds for the Troll to down you. Therefore, it is probably a good idea to run away.

5.3 (5 points). Later that night, you friend boasts about a rare bow that they found. It has the following special ability: each round, you may choose to attack twice (normally, you are only allowed to attack once per round), but both attacks suffer a -2 penalty when determining whether or not they hit (damage is unaffected).

For example, say you use this ability and get a 10 and a 12 on the two die rolls. Normally, the Attack Value for each result would be 20 and 22, respectively, but because of the -2 penalty, the final Attack Value for each attack would instead be 18 and 20, respectively.

(a) **3 points**. The bow's damage is calculated as follows: roll an 8-sided die and add 4 to the result. Your friend has an Attack rating of 9. Should they use the bow's ability while fighting the Troll?

Answer: The expected damage (given that it hits) of the bow is $\frac{1}{8}(\sum_{i=1}^{8} i) = \frac{9}{2} + 4 = 8.5$. Normally, your friend hits the Troll with probability $\frac{9}{20}$ (he needs a 12 or better). This means that if they do not use the ability, they should expect to deal $\frac{8}{20} \cdot 8.5 = \frac{34}{10} = 3.4$ damage each round to the Troll.

Now suppose your friend activates the bow's ability. This means that, on a given attack, they will hit the Troll with probability only $\frac{7}{20}$ (because of the -2 penalty). The expected damage for one attack is now $8.5 \cdot 0.35 = 2.975$. However, since your friend gets two attacks, we can apply linearity of expectation to find the expected damage in a round (both attacks), which is just $2 \cdot 2.975 = 5.95$.

Therefore, it is to your friend's advantage to use the bow's ability (even though each shot has a lower chance to hit, the second try makes it worth it!).

(b) **2 points**. Suppose your friend goes on vacation, and you get to thinking about that special ability from (a). Let y be the probability that you will hit a given monster, and let x be the expected value of the damage you will deal if you hit (note that this is ***not*** the expected damage!). For what values of x and y is it advantageous to use that special ability?

You may assume you never do 0 damage, given that you hit.

Answer: The expected damage without the ability is simply xy. The expected damage with the ability is $2x(y - \frac{2}{20}) = 2xy - \frac{x}{5}$. To see when we expect the ability to do more damage, we set

$$xy < 2xy - \frac{x}{5}.$$

We see that we can divide through by x (since it is never 0), and our inequality now only depends on y. Solving, we get

$$\begin{aligned} y &< 2y - \frac{1}{5} \\ \frac{1}{5} &< y. \end{aligned}$$

Therefore, as long as y is more than 20%, you should always use the special ability.

Question 6 (15 points) I infer that I should scream GO Bears!

The valiant California Golden Bears are in the fight of their lives against the hated Stanford Cardinal. So far, the Bears have been completely unable to get any offensive production. At halftime, star quarterback Zach Maynard has a talk with his coach, and they come to the following conclusions:
Stanford's defense can be in one of two formations, Zone coverage or Man coverage. Because they don't know which formation Stanford will be aligned in to start the second half, they assume both are equally likely.
If Stanford is in Man coverage, Zach knows that he complete a pass with probability 0.8. However, if they are in Zone coverage, he can complete a pass with probability 0.3.
[For each part, you may define any quantity from a previous part as a variable and use it with no penalty.]

6.1 (5 points). On the opening drive of the second half, Zach is able to complete a pass. What is his new guess for the probability that Stanford is in Zone coverage?
Answer: We first define the following events:

$$\begin{aligned} Z &= \text{Stanford is in Zone coverage} \\ C &= \text{Zach Maynard completes a pass} \end{aligned}$$

We are given:

$$\begin{aligned} Pr[Z] &= 0.5 \\ Pr[C|Z] &= 0.3 \\ Pr[C|\bar{Z}] &= 0.8 \end{aligned}$$

To compute $Pr[C]$, we use the total probability rule:

$$\begin{aligned} Pr[C] &= Pr[C|Z]Pr[Z] + Pr[C|\bar{Z}]Pr[\bar{Z}] \\ &= 0.3 \cdot 0.5 + 0.8 \cdot 0.5 \\ &= 0.15 + 0.4 \\ &= 0.55 = \frac{11}{20} \end{aligned}$$

Now we use Bayes's rule to find $Pr[Z|C]$:

$$\begin{aligned} Pr[Z|C] &= \frac{Pr[C|Z]Pr[Z]}{Pr[C]} \\ &= \frac{0.3 \cdot 0.5}{0.55} \\ &= \frac{0.15}{0.55} \\ &= \frac{3}{20} \cdot \frac{20}{11} \\ &= \frac{3}{11} \end{aligned}$$

6.2 (5 points). Through extensive film study, Zach Maynard knows that if the Stanford defense is in Zone coverage for a given play, they will stay in Zone coverage with probability $\frac{3}{4}$ the following play.

Given Zach's updated belief about Stanford's defense (that is, that Zach just completed a pass), what will be his guess for their formation on the following play?
You may assume (as many offensive coordinators do) that each play depends only on the previous play.
Hint: if an event A depends only on event B, then $\forall C.(Pr[A|B,C] = Pr[A|B])$.

Answer: Let X_i be a random variable denoting the i-th play's formation (Zone or Man), and let C_i be a random variable denoting the i-th passing play's success (complete or incomplete). From the previous part, we know that $Pr[X_1 = \text{Zone}|C_1 = \text{complete}] = \frac{2}{11}$ and that $Pr[X_1 = \text{Man}|C_1 = \text{complete}] = \frac{9}{11}$. To calculate $Pr[X_2|C_1]$, we simply use the total probability rule:

$$\begin{aligned}
Pr[X_2 = \text{Zone}|C_1 = \text{complete}] &= (Pr[X_2 = \text{Zone}|X_1 = \text{Zone}]) \cdot (Pr[X_1 = \text{Zone}|C_1 = \text{complete}]) \\
&\quad +(Pr[X_2 = \text{Zone}|X_1 = \text{Man}]) \cdot (Pr[X_1 = \text{Man}|C_1 = \text{complete}]) \\
&= \frac{3}{4} \cdot \frac{3}{11} + \frac{1}{4} \cdot \frac{8}{11} \\
&= \frac{17}{44} \approx 0.39 \approx \frac{1}{5}
\end{aligned}$$

This means that $Pr[X_2 = \text{Man}|C_1 = \text{complete}] = 1 - \frac{17}{44} = \frac{27}{44} \approx 0.61 \approx \frac{3}{5}$.

Note: Because we assume that a play depends only on the previous play, $Pr[X_2|X_1, C_1] = Pr[X_2|X_1]$.

Comment: There was an error in the problem statement, where the switch probability when Stanford was in Man coverage was omitted. Many students assumed that it was $\frac{1}{2}$ because in the absence of information, you cannot update your prior. We accepted this solution, as well as ones that assumed the switch probability was identical to Zone coverage (it was later announced that this was the case).

6.3 (5 points) Suppose that Zach Maynard completes another pass on the following play. Given everything that has happened, what is his guess for the probability that Stanford is in Zone coverage now?

Answer: This is just part 4.1, but using the answer from the part 4.2 as the new prior distribution:

$$
\begin{aligned}
Pr[C] &= Pr[C|Z]Pr[Z] + Pr[C|\bar{Z}]Pr[\bar{Z}] \\
&= 0.3 \cdot 0.39 + 0.8 \cdot 0.61 \\
&\approx 0.12 + 0.48 \\
&= 0.6
\end{aligned}
$$

$$
\begin{aligned}
Pr[Z|C] &= \frac{Pr[C|Z]Pr[Z]}{Pr[C]} \\
&= \frac{0.3 \cdot 0.39}{0.6} \\
&\approx \frac{0.12}{0.6} \\
&= \frac{1}{5}
\end{aligned}
$$

Question 7 (15 points)

CS70 teaching staff are trying to infer how many GSIs would need to be present at the final exam to be able to answer questions students might have on the exam. Each GSIs recorded the statistics of how often students enrolled in their sections ask questions during the last midterm exams. You can assume the students ask questions independently from each other. You can assume the number of questions students ask follows Poisson distribution.

- Hyun Oh's students ask N_1 number of questions per every 60 minutes on average.
- Victor's students ask N_2 number of questions per every 90 minutes on average.
- Dapo's students ask N_3 number of questions per every 30 minutes on average.

7.1 (7 points). What is the probability that students ask at least two questions for the three hour(180 min) final exam?

Answer: Let the poisson parameters $\lambda_1, \lambda_2, \lambda_3$ denote the average number of questions students from Hyun Oh, Victor and Dapo's section during the three hour final exam. Then $\lambda_1 = 3N_1, \lambda_2 = 2N_2, \lambda_3 = 6N_3$. Also let the random variables X_1, X_2, X_3 denote the number of questions students from the three groups ask during the final exam. The the question is asking, $\Pr[X_1 + X_2 + X_3 \geq 2]$

$$\begin{aligned}\Pr[X_1 + X_2 + X_3 \geq 2] &= 1 - \Pr[X_1 + X_2 + X_3 = 0] - \Pr[X_1 + X_2 + X_3 = 1] \\ &= 1 - \Pr[X_1 = 0, X_2 = 0, X_3 = 0] - \Pr[X_1 = 1, X_2 = 0, X_3 = 0] \\ &- \Pr[X_1 = 0, X_2 = 1, X_3 = 0] - \Pr[X_1 = 0, X_2 = 0, X_3 = 1]\end{aligned}$$

From independence,

$$\begin{aligned}\Pr[X_1 + X_2 + X_3 \geq 2] &= 1 - e^{-\lambda_1}e^{-\lambda_2}e^{-\lambda_3} - \lambda_1 e^{-\lambda_1}e^{-\lambda_2}e^{-\lambda_3} - \lambda_2 e^{-\lambda_2}e^{-\lambda_1}e^{-\lambda_3} - \lambda_3 e^{-\lambda_3}e^{-\lambda_1}e^{\lambda_2} \\ &= 1 - e^{-(\lambda_1+\lambda_2+\lambda_3)}(1 + \lambda_1 + \lambda_2 + \lambda_3)\end{aligned}$$

7.2 (8 points). If we let Z be the number of questions asked by students from one of the three sections with least number of questions, what is the probability that $Z \geq 2$?

Answer: The question is asking $\Pr[\min(X_1, X_2, X_3) \geq 2]$. Again from independence,

$$\begin{aligned}\Pr[\min(X_1, X_2, X_3) \geq 2] &= \Pr[X_1 \geq 2 \text{ and } X_2 \geq 2 \text{ and } X_3 \geq 2] \\ &= \Pr[X_1 \geq 2]\Pr[X_2 \geq 2]\Pr[X_3 \geq 2] \\ &= \left(1 - e^{-\lambda_1} - \lambda_1 e^{-\lambda_1}\right)\left(1 - e^{-\lambda_2} - \lambda_2 e^{-\lambda_2}\right)\left(1 - e^{-\lambda_3} - \lambda_3 e^{-\lambda_3}\right)\end{aligned}$$

Question 8 (15 points)

You and your best friend Ned Stark are watching your mutual enemy Jaime Lannister throw darts at a board. It is known that each time Jaime tosses a dart at the dartboard, the distances R_i from the bullseye can be modeled as a Uniform$(0, 1)$ random variable (the distance is uniform on (0,1), and the radius of the dartboard is not necessarily 1.)

Moreover, it is known that the tosses are independent.

You and Ned agree to the following wager

1. You will both observe Jaime toss 100 darts at a dart board, each a distance R_i from the bullseye.
2. If the sum of the distances is greater than or equal to 55, Ned wins \$10.
3. If the sum is between $50 - \delta$ and $50 + \delta$, you win \$1.
4. Otherwise, no one wins the wager.

Let X denote the sum of the the distances of the 100 dart tosses.

8.1 (1 points). Compute $\mu = E[X]$. **Answer:** 50

8.2 (2 points). Compute $V = Var[X]$. **Answer:** 100/12

8.3 (2 points). Using Markov's inequality, approximate the probability that Ned wins. Your answer should be expressed in terms of some combination of μ, V (note that your final answer might or might not use all of these variables.)

Answer:

$$\Pr(\text{Ned Wins}) = \Pr(\sum_{i=1}^{100} R_i \geq 55) \leq \frac{\mu}{55}.$$

8.4 (2 points). Using Chebyshev's inequality, approximate the probability that Ned wins. Your answer should be expressed in terms of some combination of μ, V (note that your final answer might or might not use all of these variables.)

Answer:

$$\Pr(\text{Ned Wins}) = \Pr(\sum_{i=1}^{100} R_i - \mu \geq 55 - \mu) \leq \frac{V}{[55-\mu]^2} = 1/3.$$

8.5 (3 points). Let the function $\Phi(x) = \Pr(Z \leq x)$, where Z is a $N(0,1)$ random variable. Approximate the probability that Ned wins using the central limit theorem (your final answer should involve $\Phi(a)$ for some number a.) Your answer should be expressed in terms of some combination of μ, V (note that your final answer might or might not use all of these variables.)

Answer: CLT says that

$$X \approx N(\mu, \sigma^2).$$

So

$$\begin{aligned}\Pr(\text{Ned Wins}) = \Pr(X - 50 \geq 5) &= \Pr(\frac{X-\mu}{\sigma} \geq \frac{5}{\sigma}) \\ &\approx \Pr(Z \geq \frac{5}{\sigma}) = 1 - \Phi(\frac{5}{\sigma}).\end{aligned}$$

8.6 (5 points). For the game to be fair, your expected payoff and that of Ned should be equal. We will now calculate what δ must be for the game to be fair. In other words, if θ_n denotes the probability of Ned winning and θ_u denotes your probability of winning, then $10\theta_n = \theta_u$.

Assume that θ_n is 0.05 (this may or may not be the answer to previous parts of the problem, don't necessarily assume that it is, or use those previous calculated values of θ_n for this part.)

Using the central limit theorem, (approximately) choose δ so that the game is fair.

You will find it useful to use $\Phi^{-1}(x)$ to denote the inverse of the function $\Phi(x)$. In other words, given any x, if $y = \Phi(x)$, then $\Phi^{-1}(y) = x$.
Answer: We want

$$10\,\Pr(\text{Ned Wins}) = 0.5 = \Pr(\text{I win}) = \Pr(|X - 50| < \delta).$$

CLT says that

$$X \approx N(\mu, \sigma^2).$$

So

$$\begin{aligned}
\Pr(|X - \mu| < \delta) &\approx \Pr(|N(0, \sigma^2)| < \delta) \\
&= 1 - 2\,\Pr(N(0, \sigma^2) > \delta) = 0.5 \\
&\implies \Pr(N(0, \sigma^2) > \delta) = 0.25 \\
&\implies \Pr(Z < \frac{\delta}{\sigma}) = 0.75 \\
&\implies \Phi(\frac{\delta}{\sigma}) = 0.75 \\
\implies \frac{\delta}{\sigma} = \Phi^{-1}(0.75) &\implies \delta = \sigma\Phi^{-1}(0.75)
\end{aligned}$$