Name: ______________________________ Exam 01 (Midterm Part 2 Take Home, Open Everything)

To help you budget your time, questions are marked with *s. One * indicates a straightforward question testing foundational knowledge. Two ** indicate a more challenging question requiring use of several key concepts. Three *** indicate a challenging question requiring a clever use of key concepts and/or techniques not covered in class yet. Four **** indicates serious fun. **Show your work on all problems.**

If you can't solve a question completely, set it up and solve it as far as you can. If you require a laptop to get your final solution, describe in detail what you would do to solve the problem.

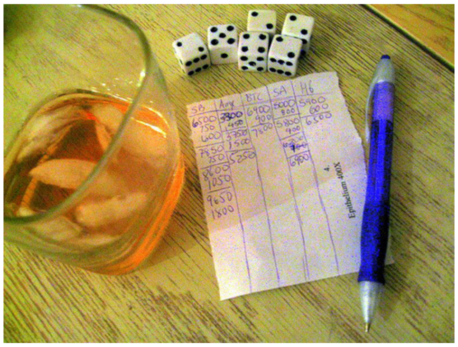

image used without permission from http://www.farklerules.com/wp-content/uploads/2008/08/flickrfarkle.jpg

Basic Farkle Scoring Rules:
Each 1 = 100 pts
Each 5 = 50 pts
Three 1's = 1000 pts
Three 2's = 200 pts
Three 3's = 300 pts
Three 4's = 400 pts
Three 5's = 500 pts
Three 6's = 600 pts
Straight (1-2-3-4-5-6) = 1000 pts

Game Overview:
The players take turns rolling the dice, with the objective of having the highest score above 10,000 in the final round.

During each player's turn, they initially roll six dice trying to score points. As long as they score at least one point, they can either bank their points and pass the dice, or remove the scoring dice from play and roll the remaining dice.

If the dice you roll do not score any points, you pass the dice and you Farkle, losing all points accumulated for that turn.

If the player manages to score on all six dice, they have "hot dice" and may choose to roll all six dice again, or they can bank the points and pass the dice.
At the end of the player's turn, they write down any points scored and pass the dice clockwise.

***1 - 5pts)** If rolling just one die, what is the probability that you will Farkle on that roll, i.e. **P(Farkle | rolling 1 die)**? Write your solution as a decimal accurate to four places.

P(roll 1,2,3, or 4) = 4/6 = 2/3 = 0.6667

***2 - 5pts)** If rolling just two dice, what is the probability that you will Farkle on that roll, i.e. **P(Farkle | rolling 2 dice)**? Write your solution as a decimal accurate to four places.

P(roll 1,2,3, or 4 on both dice) = (4/6)^2 = 4/9 = 0.4444

http://www.mobiletopsoft.com/pocket-pc/newimg/farkle_ppc_straight.gif

****3 - 10pts)** If rolling all six dice, what is the probability that you will roll a straight on that roll, i.e.
P(Straight | rolling 6 dice)?
Write your solution as a decimal accurate to four places.

P(roll 1-2-3-4-5-6 in any order) = 6! * 1/6^6
= 720 / 46,656
= 0.0154

*****4 - 15pts)** Suppose that in a game of Farkle, there will be 3 times more rolls that use all six dice than rolls that use only five dice. If a single roll scores 1200 points, what is the probability the roll used six dice? That is, what is **P(rolled six dice | 1200 points on that roll)**?

Hint:

P(1200 points on that roll | rolled six dice) = 1/6^6 + 6*1/6^5*4/6 + 20*1/6^6 = 0.0009645062

P(1200 points on that roll | rolled five dice) = 1/6^5 = 0.0001286008

P(1200 points on that roll | rolled less than five dice) = 0

$$P(\text{rolled 6d} \mid \text{1200 pts}) = \frac{P(\text{1200 pts} \mid \text{rolled 6d})*P(\text{rolled 6d})}{P(\text{1200 pts} \mid \text{rolled 6d})*P(\text{rolled 6d}) + P(\text{1200 pts} \mid \text{rolled 5d})*P(\text{rolled 5d})}$$

$$= \frac{0.0009645062*P(\text{rolled 6d})}{0.0009645062*P(\text{rolled 6d}) + 0.0001286008*P(\text{rolled 5d})}$$

$$= \frac{0.0009645062*(\ 3*P(\text{rolled 5d})\ )}{0.0009645062*(\ 3*P(\text{rolled 5d})\ ) + 0.0001286008*P(\text{rolled 5d})}$$

$$= \frac{3*0.0009645062}{3*0.0009645062 + 0.0001286008}$$

$$= 0.9574$$

http://depositphotos.com/5541834/stock-photo-Fun-and-Games---Words-on-Three-Red-Dice.html

***5 - 10pts)** Cautious Carl's strategy is to take all points available on any roll and to bank his points whenever he has three dice or less left. I've taken a sample of 25 turns using Carl's strategy and found his strategy scored an average of 468 points per turn with a sample standard deviation of 412.2802. **Using your method for normally distributed data with an unknown variance, create a 95% confidence interval for the true mean number of points per turn using Carl's strategy.** Write your solution as (LB, UB) accurate to two decimal places.

My data: 600, 2000, 0, 700, 450, 550, 250, 200, 250, 550, 200, 200, 450, 350, 300, 350, 600, 0, 1100, 200, 250, 350, 400, 1000, 400

$x_{bar} \pm t_{24,0.975} * s / \text{sqrt}(n)$

468 ± 2.064 * 412.2802 / sqrt(25)

468 ± 2.064 * 412.2802 / sqrt(25)

468 ± 170.1893

(297.81, 638.19)

In R:

data = c(600, 2000, 0, 700, 450, 550, 250, 200, 250, 550, 200, 200, 450, 350, 300, 350, 600, 0, 1100, 200, 250, 350, 400, 1000, 400)

t.test(data)$conf.int

(297.8191, 638.1809)

Note the disagreement in second decimal due to my rounding the standard deviation:
print(sd(data), digits=16)
412.2802444939607

http://depositphotos.com/4434534/stock-photo-Take-a-Risk---Roll-the-Dice.html

***6 - 12pts)** Gambling Garth believes a Farkle strategy is worthless unless it averages 500 points per turn. Garth takes all points available on a roll and keeps rolling on any number of dice until he has 400 points or more, then he banks if he has only three dice or less left to roll. I've taken a sample of 30 turns using Garth's strategy and found his strategy scored an average of 415 points per turn with a sample standard deviation of 437.4988. **Using your method for normally distributed data with an unknown variance, perform a two-sided hypothesis test of whether Garth's strategy averages 500 points per turn.** Fill out all the elements indicated below.

My data: 750, 1050, 500, 400, 0, 1150, 0, 0, 0, 0, 0, 1550, 400, 0, 400, 0, 700, 550, 600, 450, 600, 500, 400, 0, 0, 0, 0, 400, 1350, 700

Ho: $\mu = 500$

Ha: $\mu \neq 500$

Test Statistic (formula): TS = (xbar - 500) / (s / sqrt(n))

Test Statistic (observed value): TS = (415 - 500) / (437.4988 / sqrt(30))

= -85 / 79.87599

= -1.06415

Rejection Region at α = 0.05: | TS | = 2.045

P-value (in class you may use the table to give bounds for p-value, i.e. lb < p-value < ub):

0.20 < p-value < 0.40

Conclusion at α = 0.05: There is not sufficient evidence at a 5% significance level to reject the possibility that the true mean points per turn for Garth's strategy equals 500. A 95% CI for the mean points per turn is (251.64, 578.36).

R solution: data = c(750, 1050, 500, 400, 0, 1150, 0, 0, 0, 0, 0, 1550, 400, 0, 400, 0, 700, 550, 600, 450, 600, 500, 400, 0, 0, 0, 0, 400, 1350, 700)
t.test(data, mu=500)

****7 - 13pts)** Cautious Carl and Gambling Garth are arguing over whose strategy is better. Garth says all those turns with scores less than 400 don't help Carl. Carl says Garth's method yields just as many scores under 400, but they are all coming from farkles in Garth's method. I observed that out of 25 turns, Carl had 13 scores under 400. Out of 30 turns, Garth had 12 scores under 400. **Using your asymptotically normal test that pools the data to estimate the standard error, perform a two-sided hypothesis test of whether Carl's claim that both methods have the same proportion of scores less than 400.** Fill out all the elements indicated below.

Ho: $\theta_{Carl} - \theta_{Garth} = 0$

Ha: $\theta_{Carl} - \theta_{Garth} \neq 0$

Test Statistic (formula): $\theta_{pooled} = (N_{Garth}*\theta hat_{Garth} + N_{Carl}*\theta hat_{Carl}) / (N_{Carl}+N_{Garth})$

$TS = (\theta hat_{Carl} - \theta hat_{Garth}) / sqrt(\theta_{pooled}*(1- \theta_{pooled})*(1/N_{Carl}+1/N_{Garth}))$

Test Statistic (observed value): $\theta_{pooled} = (13+12) / (25+30) = 25/55 = 5/11$

TS = (13/25 - 12/30) / sqrt((5/11)*(6/11)*(1/25+1/30))

= 0.8899

Rejection Region at α = 0.10: $|TS| > Z_{0.95} = 1.645$

P-value: 2*(1-0.8133) = 2*0.1867 = 0.3734

Conclusion at α = 0.10: There is not sufficient evidence at a 10% significance level to reject the possibility that the true proportions for scoring less than 400 pts in a turn are the same for Carl's and Garth's strategies.

R solution: chisq.test(matrix(c(13,12,12,18),nrow=2), correct=F)

*****8 - 10pts)** In my sample of Garth's method, it took 23 turns to reach 10,000 points. The number of turns to reach 10,000 is a great example of a Poisson random variable. Using my one sample (23 turns) and the Poisson distribution, **create an exact 95% confidence interval for λ**, the number of turns to reach 10,000 points using Garth's strategy.

X ~ Poisson(λ).

For a 95% CI of (λ_{lb}, λ_{ub}), we have
P(X $\geq$ 23 | λ_{lb}) = 0.025 and P(X $\leq$ 23 | λ_{ub}) = 0.025

Solve 1 - P(X $\leq$ 22 | λ_{lb}) = 0.025
1 - ppois(q=22, lambda=14.580026)
= 0.02499998
1 - ppois(q=22, lambda=14.580027)
= 0.025
1 - ppois(q=22, lambda=14.580028)
= 0.02500002
λ_{lb} = 14.580027

Solve P(X $\leq$ 23 | λ_{ub}) = 0.025
ppois(q=23, lambda=34.51130)
= 0.02499993
ppois(q=23, lambda=34.51129)
= 0.02500003
λ_{ub} = 34.5113

(14.5800, 34.5113)

Aside: The above solution is a really useful way to understand exact confidence intervals. It's the right way to think through the problem to help you get the key idea: solve for a λ_{lb} that makes X=23 seem big and a λ_{ub} that makes X=23 seem small. If you're interested in how R or Stata solve the problem, it turns out there is a relationship between the Poisson and Chi-squared distributions that make this problem solvable in closed form.
(λ_{lb}, λ_{ub}) = c(qchisq(0.025, 2*x)/2, qchisq(0.975, 2*(x+1))/2)
= c(qchisq(0.025, 2*23)/2, qchisq(0.975, 2*(23+1))/2)
= (14.58003, 34.51129)
= (14.5800, 34.5113)

*****9 - 10pts)** Both Garth and Carl took all points available on every roll. I'm not convinced this is always the best strategy. Suppose you just rolled five dice and rolled a 1 - 2 - 2 - 3 - 5 and you want to roll again. **Should you keep the 1 and the 5 and roll three dice or keep the 1 only and roll four dice?** Justify your answer.

A good in-class solution:

This will take some complex computation, but let me set up the solution and explain the logic behind it.

Let X_{keep1} = the total bankable points after keeping the 1 and rolling the remaining 4 dice.

$E[X_{keep1}]$ = 100*(1-P(Farkle rolling 4 dice)) + E[points scored on a single roll of 4 dice]

Let X_{keep2} = the total bankable points after keeping the 1 & 5 and rolling the remaining 3 dice.

$E[X_{keep2}]$ = 150*(1-P(Farkle rolling 3 dice)) + E[points scored on a single roll of 3 dice].

Logic: The strategy with the higher expectation is better, so I need to calculate and compare $E[X_{keep1}]$ and $E[X_{keep2}]$. This is a simplification in that it stops after 1 roll. So the case for X_{keep1} when it rolls only a scoring 1 (100pts) on the next roll will be undervalued because in that case the roller would be able to roll the remaining three dice to score more points. In other words, they'd be in the keep2 position with 200 points on hand instead of the 150. Thus, the comparison I propose isn't completely fair. A modest improvement would be to define X_{keep1} as the total bankable points after keeping the 1 and rolling the remaining 4 dice AND rolling again with the remaining three dice if they only scored on one of the 4 dice they just rolled.

My intuition suggests the benefit from rolling 4 dice in terms of P(scoring a triplet) will outweigh sacrificing the 50 points, but I'll have to do the math to find out.

******10 - 10pts)** A websearch for *farkle strategy* yields many pages of advice. One page suggests the following strategy. If you are considering throwing...

6 dice: Just do it! Don't worry about it.
5 dice: Stop if you already have 2000 points or more. Otherwise go ahead and throw.
4 dice: Stop if you already have 1000 points or more. Otherwise go ahead and throw.
3 dice: Stop if you already have 500 points or more. Otherwise go ahead and throw.
2 dice: Stop if you already have 400 points or more. Otherwise go ahead and throw.
1 dice: Stop if you already have 300 points or more. Otherwise go ahead and throw.
source: http://EzineArticles.com/4024070

This approach can be generalized as follows. Specify six stopping thresholds (δ_6, δ_5, δ_4, δ_3, δ_2, δ_1), one threshold for each of the possible number of dice you're about to throw. The strategy above can be written as (δ_6, δ_5, δ_4, δ_3, δ_2, δ_1) = (∞, 2000, 1000, 500, 400, 300). Come up with a better strategy, that is a better set of (δ_6, δ_5, δ_4, δ_3, δ_2, δ_1), and **prove your strategy is better.** Include all the details of your proof, e.g. mathematical calculations, data collection methods, computer programs, etc.

A good in-class solution:

This problem will take some advanced R coding to solve, but I can set up the solution and explain the logic behind it.

I need to find a set of δ such that my strategy will reach 10,000 points in fewer turns than the Ezine strategy does, i.e. I reach 10,000 points first. No strategy can guarantee this, so I need:
P(my strategy reaches 10,000 points before the Ezine strategy does) > 0.5

I could try to simplify by solving for the expected points scored per turn or the expected number of turns to reach 10,000. Both of those would be good proxies, but since I'm going to write a simulation to solve this, it is better to avoid the proxies and specifically solve for the probability above. I'll sketch out the R code below.

The hard part will be writing a function that plays Farkle using a given set of δs. Let turns() be a function that returns the number of turns required to hit 10,000 in one randomly simulated game of Farkle.

Let win = 1 if turns(my δs) < turns(Ezine δs)
win = 0 if turns(my δs) > turns(Ezine δs)
win = 0.5 if turns(my δs) = turns(Ezine δs) -- Note, this captures the 50-50 chance of winning if who goes first is determined by a fair coin toss.

Let wins = win from a large number of simulated games of Farkle, e.g. 200,000 games.

mean(wins) = P(my strategy reaches 10,000 points before the Ezine strategy does).