APMA 1650 – FINAL EXAM – STUDY GUIDE

This is the study guide for the exam, and is just a **guide** to help you study. Think of it as a course summary rather than "this is exactly the questions I'm going to ask you on the exam."

Format: There are about 8-10 questions on the exam, all of them free response, no multiple choice.

You're allowed one 2-sided 8.5×11 cheat sheet; no other formulas will be provided. If necessary, I will give you either a $Z-$table or "fake" $Z-$values, like on the second midterm or the final review, same for $t-$tables or Chi-Square tables.

1. Probability Basics (Lectures 1-2)

- Write down the sample space of an experiment
- Write down an event in set notation
- Define: union, intersection, and complement
- Know the distributive laws and De Morgan's laws
- State the probability axioms
- Give an example of three sets A, B, C such that $A \cap B \cap C = \emptyset$ but A, B, C are not pairwise disjoint
- Show that for the unif probability distribution, we have $p = \frac{1}{n}$

2. Probability and Counting (Lectures 3-4)

- Use the mn rule to count the total number of possibilities. Use this when you're counting two independent things and the order matters, like $(1,6) \neq (6,1)$
- Find the number of permutations of r objects from n objects. Use this when the order matters, like rearranging letters `PIE` or license plate numbers
- Find the number of combinations of r objects from n objects. Use this when the order doesn't matter, like `123 = 321`. There are many problems that used this method:
    - Raffle tickets
    - Undergraduate/Graduate Committee
    - Poker
    - Antenna
- Use the box and star diagram to count the different possibilities of, say, donuts in a store. Use this when the order doesn't matter and when repetition is allowed (like 4 blueberry donuts)
- State the binomial theorem
- Solve a problem using multinomial coefficients, like putting 15 students in 3 different groups of 5. It's totally ok to use binomial coefficients instead, but in that case make sure to simplify your answer to look something like $\frac{15!}{5!5!5!}$
- Study the `MISSISSIPPI` problem

3. Conditional Probability, Independence, and Bayes (Lectures 5-7)

- Define: $P(A|B)$. It's useful in practice to know both versions, the one with "Probability of A given B" and the $\frac{P(A\cap B)}{P(B)}$ version.
- Show that two events are independent, using $P(A|B) = P(A)$ or $P(A \cap B) = P(A)P(B)$
- Use the multiplicative law
- Define: Additive Law
- Use a tree diagram to solve a problem. This is useful when you have several cases and the events are not independent
- State Bayes' Theorem, but you don't need to prove it. This is useful if you want to calculate $P(A|B)$ but you know $P(B|A)$
- Know the law of total probability and the general Bayes' Theorem. Those are useful when you can partition your universe into disjoint events/classes
- There are many problems that use this method, like
  - Urn Problem
  - Card Game Problem
  - HIV Case Study, I would provide you $P(H) = 0.00376$

4. Random Variables (Lectures 8–9)

- Define: Random Variable, $(Y = y)$, probability mass function $p(y)$, sample space induced by a random variable

- Find $p(y)$ either as a table or as a histogram
- Find the expected value of a discrete random variable
- Use the method of indicators to find expected values, like the concierge problem in lecture or the marble problem on the homework or the ice cream problem on the first midterm. Use this method when you have n people and when it is hard to calculate $P(X = i)$ directly.
- Find $E(g(X))$ where g is a real valued function
- Find the variance of a discrete random variable
- Know the magic variance formula and know how to derive it
- In general we don't have $\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y)$ but it is true if the random variables are independent

5. Distributions (Lectures 10–13)

- Define: Bernoulli trial, Bernoulli Random Variable.
- You use a Bernoulli random variable is if you have **one** experiment, like **one** coin toss with two outcomes, success and failure
- Find the expectation and var of a Bernoulli Random Variable
- Define: Binomial Distribution and know how to find $p(y)$
- You use a binomial distribution if you have a **sequence** of Bernoulli trials and you count the number of successes.
- Find the expectation, variance of a Binomial Random Variable and show that the sum of the probabilities is 1

- Define: Geometric Distribution and know how to find $p(y)$
- You use the geometric distribution if you want to know how long it takes to get your first success
- Know the formula of the geometric series
- Find the expectation of a Geometric Random Variable and show that the sum of probabilities is 1. You don't need to prove the formula for the variance
- Know the memoriless property of the geometric distribution, that is, the geometric distribution "forgets" about the first m trials, as if we started from scratch
- Define: Hypergeometric distribution and know how to find $p(y)$
- You use the hypergeometric distribution to count the number of successes when there is no replacement. Compare this with the binomial distribution where there is replacement
- Define: Poisson Distribution, no need to know how to find $p(y)$
- No need to know how to construct the Poisson distribution
- Find the expectation of a Poisson Random Variable and show that the sum of probabilities is 1. You don't need to prove the formula for the variance
- You use the Poisson distribution to count the number of events that occur during a fixed time interval

6. Continuous Random Variables (Lectures 13-14)

Given a function $f(y)$

- Find the value of c for which $f(y)$ is a pdf
- Find $P(a \leq Y \leq b)$
- Find the cdf $F(y)$ and use F to calculate $P(a \leq Y \leq b)$. The normal distribution is an excellent example
- Find the first/third quartiles and median of a random variable
- Find $E(Y)$ and $\text{Var}(Y)$

7. Continuous Distributions (Lectures 14-15 + 17)

- Solve problems using the uniform distribution on $[a, b]$
- Find the $E(Y)$ and $\text{Var}(Y)$ where $Y \sim$ Unif (a, b)
- Use the $Z-$table to find $P(Z \leq a), P(Z \geq a), P(a \leq Z \leq b)$
- Sometimes we write $Y \sim N(\mu, \sigma)$ or $Y \sim N(\mu, \sigma^2)$. In both cases it is implied that the variance is σ^2
- Know the $68 - 95 - 99.7$ rule. Beware if you see those numbers, it's likely that you use that rule in this case.
- The transformation $Z = \frac{Y-\mu}{\sigma}$ is **super** important, it allows you to go from $Y \sim N(\mu, \sigma^2)$ to $Z \sim N(0, 1)$ and vice-versa
- Know the pdf of the exponential distribution. It measures the amount of time between two subsequent events

- Show that the exponential distribution integrates to 1 and prove the formulas for $E(Y)$, $\text{Var}(Y)$, $F(y)$ and the median
- Prove the memoriless property of the exponential distribution

8. Markov and Chebyshev's Inequality (Lectures 17-18)

- State and prove Markov's and Chebyshev's inequalities
- You use Markov if you don't know *anything* about Y except for its mean μ
- You use Chebyshev if you don't know *anything* about Y except for its mean μ and variance σ^2
- If you know the distribution of Y (like normal or exponential) use that distribution instead! Why use an approximate answer if you know the exact one ☺
- Prove that fact about "deviating at least k standard deviations"

9. Multivariate Distributions – discrete (Lectures 18-19)

- Find the joint distribution function $p(y_1, y_2)$ of Y_1 and Y_2
- Show that $p(y_1, y_2)$ is a valid joint distribution function
- Find the marginal distributions $p_1(y_1)$ and $p_2(y_2)$
- Find the expected values $E(Y_1)$ and $E(Y_2)$
- Find the conditional distribution $p(y_1|y_2)$
- Show that Y_1 and Y_2 are independent

10. Multivariate Distributions – continuous (Lec 19-22)

Given a function $f(y_1, y_2)$

- Find the value of c for which f is a valid joint density function
- Calculate double integrals. A picture is absolutely crucial here.
- Find $P((Y_1, Y_2) \in A)$ where A is a region in the plane
- Find the marginal densities $f_1(y_1)$ and $f_2(y_2)$
- Find the expected values $E(Y_1)$ and $E(Y_2)$
- Find the conditional density $f(y_1|y_2)$
- Find the conditional expected value $E[Y_1|Y_2 = y_2]$
- Show Y_1 and Y_2 are independent
- Know the density of the uniform distribution
- Find $E[g(Y_1, Y_2)]$ both in the discrete case and the continuous case
- Show that if Y_1 and Y_2 are independent then $E(Y_1Y_2) = E(Y_1)E(Y_2)$
- Find $\text{Cov}(Y_1, Y_2)$ and ρ
- Prove the Magic Covariance Formula
- Show that if Y_1 and Y_2 are independent, then $\text{Cov}(Y_1, Y_2) = 0$
- Know the general formula for $\text{Var}(Y_1 + Y_2)$ no need to prove it

11. Sampling Distributions (Lectures 23-26)

- Skip the introduction but know the definition of $\bar{Y}$ and S^2
- Find $E(\bar{Y})$ and $\text{Var}(\bar{Y})$
- In general the distribution of $\bar{Y}$ is unknown except in the case where the $Y_i \sim N(\mu, \sigma^2)$
- Know the transformation $Z = \frac{\overline{Y}-\mu}{\sigma/\sqrt{n}}$
  This only works if the Y_i are normal!
- Know the definition of the chi-square distribution and how to use the chi-square table
- The chi square distribution is useful because **IF** the Y_i are normal because then $\left(\frac{n-1}{\sigma^2}\right) S^2$ has a chi-square distribution
- Use the method of splitting the difference
- Know the definition of the Student's $t-$ distribution and how to use the $t-$table
- The $t-$distribution is useful **IF** all the Y_i are normal and we don't know σ. If you know σ you would use the normal distribution instead!
- Know the Central Limit Theorem. This is useful if you don't know the distribution of Y_i and the sample size is large

12. Estimators (Lecture 27)

- Define: Estimator, Point Estimator, Interval Estimator
- Show that $\hat{\theta}$ is biased or unbiased, like $\overline{Y}$ and $\hat{p} = \frac{Y}{n}$

- Find $\text{MSE}(\hat{\theta})$ and prove the Magic MSE formula
- Look at the two examples about the difference in populations. You don't need to memorize the table
- Understand why we use $n-1$ in the sample variance S^2. You don't need to memorize the full proof

13. Confidence Intervals (Lectures 28-29)

- Construct a confidence interval if the sample size is large, this boils down to CLT + convert to Z
- This includes the case with differences, in that case you just use that $\text{Var}(\hat{p_1} - \hat{p_2}) = \text{Var}(\hat{p_1}) + \text{Var}(\hat{p_2})$
- Construct confidence interval if the sample size is small. Need to assume that the Y_i are normal + convert to Student's t
- Ignore the last example with the pooled variance estimator S_p^2

14. Consistency and MLE (Lectures 32-33)

- Define: $\hat{\theta}_n$ is consistent for θ
- Show $\hat{\theta}_n$ is consistent for θ. I'll only ask about the unbiased case, where you use the Magic Consistency formula, like for $\overline{Y_n}$
- Prove the Magic Consistency Formula
- Find the Method of Moments Estimator
- Ignore the motivation for MLE

- Define: MLE
- Find the MLE and study the three examples from lecture

15. Hypothesis Testing (Lectures 34-36)

- Set up the 5 components of a hypothesis test (parameter of interest, alt hypothesis, null hypothesis, test statistic, RR)
- Define: α and β
- Find the rejection region given the type I error α, in the case of large samples. This boils down to CLT + convert to Z
- Need to do this for one-tailed and two-tailed hypothesis tests
- Given α and k and a value of θ_a, find the type II error β. This boils down to convert to Z but using $\theta = \theta_a$ this time
- Given α and β, find k and n. This boils down to converting to Z and solving the equations for n and k
- Define: $p-$value of a test
- Find the p-value of a test
- **Note:** All the formulas are for the case $\left\{\hat{\theta} \geq k\right\}$ For the case $\left\{\hat{\theta} \leq k\right\}$ I recommend re-deriving everything from scratch, like in the final exam review session
- **Note:** The final will **NOT** cover small sample hypothesis testing, power of tests, and likelihood ratio tests.