

1 Me

Robin Pemantle, Prof. of Math (and Stat and CIS)

Office: 4N22 DRL, 898-5978, office hours: any time, by appointment.

TA: Sneha Subramanian (ssneha@math.upenn.edu), DRL 4C9, x8-6217, office hours: Tue 4PM-6PM

Course website: Canvas.

2 Course

Stat 530-531 / Math 546-547 – graduate level introduction to probability theory. Graduate level means “with proofs, based on theory of Lebesgue measure and integration”. It does not mean “without intuition” nor “purely formal”. Compare to Stat 510. There they use Ross. He does not even prove our first result (Strong Law of Large Numbers) except in finite variance case. Furthermore, even if you don’t care about proofs, he does not even state many (most!) critical results on which classical statistical tests are based, e.g., LIL, large deviations, any theorem or criterion for Poisson approximation to be valid, etc.

Students sometimes need some help deciding whether this is the right course for them. I’ve addressed that to some degree on the course website. What it says there is to have a look at the textual materials and the first homework (already posted) and see whether you can envision being comfortable with that material within a few weeks. To the advice there, I would add: facility with writing mathematical proofs at the level of math 360-361 is going to be much more important than any specific mathematical or statistical knowledge. Mathematical analysis at the level of Penn’s Math 508-509 is recommended and if you’ve had measure theory you’ll be glad of it, but many students have not had this and in principle an undergraduate level analysis course such as Penn’s Math 360-361 should be sufficient. Most students have had undergraduate probability but those who have not usually do fine. Students who have not taken analysis (Penn 360-361) at any level usually struggle mightily.

I’ve taken care to design the course so that (1) you don’t need any prior knowledge of measure theory to take the course and (2) we won’t need to spend all that much time doing

measure theory in the course. In particular, the level of measure-theoretic technicalities we will have to deal with will be consistent with the level that enhances intuition and is necessary for learning. As you will see, in the first semester this means the word “measurable” pretty much doesn’t have to appear at all, though we do cover all the basics in order to forestall any confusion. In the second semester, we will need to go back and take care of some more of the measure theory.

The course is cross-listed between Math, Applied Math and Statistics, and there are usually students as well from Engineering, Wharton and SAS. I will try to cover the needs of this broad audience! There may be too many proofs for the engineers and non-Statistics Wharton folks, too few for the math folks, too many statistical applications for the SAS folks and too great a scope for the Statistics folks. As I get to know who you are, I’ll try to fine-tune things to address your needs.

3 Book

Durrett, “Probability: Theory and Examples” is required. My lectures often parallel this book - I would say 50% of the material essentially follows Durrett. Do not rent it or make a cheap copy of it: one thing you get out of this course is a copy of Durrett’s textbook on your shelf and the ability to use it.

In addition, I will provide a complete set of lecture notes. You can download these from Canvas. Those who do not wish to take notes can assume that any material covered in the course will be in the lecture notes. This includes that material in Durrett. Because the book by Durrett is difficult, I recommend several other texts for those who find a different voice more suitable to their style. These include: Billingsley’s “Probability and Measure”; Rosenthal’s “A First Look at Rigorous Probability Theory” and Folland’s “Real Analysis”.

4 Work

There will be thirteen homework sets and a final exam. The grading scale for the course is: everyone who is not lost should get a B+ or better. Grades below that are used mostly to signal to varying degrees that you didn’t really learn the material. Verbal descriptions of the grades are: A = Mastery of material; A- = Mastery of some of it, adequate knowledge of

most of it; B+ = Can understand most of the material line by line but has trouble putting it together to solve harder problems; B = Superficial understanding; C = Clueless; F = stopped doing any work.

Working with other students on the homework is permitted, though care should be taken not to collaborate until you have given significant thought to a problem. In any case, each student must write their own solution: either submit it in your own handwriting, or type it into a word processor yourself. The use of worked solutions written up by the TA or anyone else outside of this year's class is not permitted, even to look at, though of course you can look at these *after* you hand in your homework.

The homeworks will be graded by the grader, who is always a Math or Stat graduate student. This year it is Sneha Subramanian, whose contact info is given above. The homework is divided into weekly problem sets of 3–4 problems each. These are all already posted on the website. We will try to publish solutions immediately after the assignment is due. Feedback from prior years tells us that students very much appreciate seeing the solutions while the problems are fresh in their minds. Consequently, we cannot accept any late homework (but you are welcome to do it early – it is all there on the web).

The final exam will be worth enough to assure me that students have independent mastery of the material (in other words, the entire course grade cannot rest on collaborative work), but not so much that a student's grade will undergo radical change if the student has been keeping up with the concepts. Probably the final will count for roughly three problem sets as long as it shows that the student was not completely dependent on others for their homework performance. Please see the posted final exam guidelines for details as to the format of the final exam.