

STATS 721: Information about the Exam

The exam is three hours long. You will need a calculator. The final mark for the course is:

EITHER 10% assignments 1 to 5; 8% assignments A & B; 7% test; 75% exam,

OR if you qualify for plussage, 8% assignments A & B; 92% exam, if this mark is better.

Remember that you need **70% on 325 coursework** (assignments 1 to 5 and test) to be considered for plussage, and you need **50% on the final exam** to pass the course.

Differences between 721 and 325 exams

There are seven questions on the exam, of which two differ from the 325 exam (approximately 25%). These two questions will be on the material of Assignments A and B. **2011 is the first year that this material has been tested on the exam.** For this reason, it is not particularly helpful to look up past 721 exams, because the extra material on the past exams was different. Instead, I recommend the following:

- For martingales, revise Assignment B thoroughly. Remember that it is quite difficult to construct martingales that satisfy the optional stopping theorem, so there aren't all that many questions that can be asked! If you understand Assignment B thoroughly, I think you will find the martingale question quite easy.
- For second-order difference equations, revise Assignment A. It might also be helpful to practise on past 325 exam questions, where the characteristic equation and general solution is given. Instead of just following the 325 instructions, solve the whole equations independently as if they were 721 questions. As long as you know the steps, the procedure for solving the equations is not difficult: it just requires a sustained / independent calculation. The questions will involve second-difference equations that are quite easy to set up from first-step analysis (*not* difficult examples like Q2&3 on Assignment A) — this is pretty obvious because the questions will allot a large number of marks to solving the equations, so it would be pointless to set questions where there is a risk that people might not get the equation right to start with.

Exam Format

With the exceptions above, the format of the exams is the same as the format of the exams for 2006 to 2010, given out in class or available from www.stat.auckland.ac.nz/~stats325/exams.php. Most of the content in the exams from 2003 to 2005 is also still relevant, with a few changes. All answers to past exams are available on the website.

You should attempt **all** questions on the exam. Marks for each question part are shown on the exam paper. An Attachment with useful formulae is provided. The Attachment is exactly the same as the Attachment sections provided on the previous exams. You should read over the Attachment carefully to be sure of what is there.

The best possible preparation for the exam is **practice**. You should attempt as many of the past exams as you can. There are very few questions for which learning notes is helpful (or even possible), and you will find these anyway on the past papers, e.g. Exam 2008, Q. 2(a); Exam 2006, Q. 8(e). Start practising exam questions from day one, and just have your notes nearby to refer to as you need them. If you spend time reading your notes instead of practising questions, you will be at a disadvantage to other students.

Question difficulty

The past exams have 'difficulty' gradings next to the marks: *E* (easy), *M* (moderate), *H* (hard). These are my best guesses, so don't worry if you disagree! The other students are more likely to agree with you than with me, and this comes out in the marking. Here are some guidelines:

1. **Easy questions.** Everybody should be able to do these. If a question marked *E* seems quite substantial, it probably means it is *predictable*. Watch out for predictable questions: anything that has appeared on all of the past exams is very likely to appear on your exam too.
2. **Moderate questions.** Half of the paper is in this category. Some *M* questions border on *H*, others border on *E*, and most are in-between. You can pass the paper if you can get all of the *E* marks and a bit over half of the *M* marks.
3. **Hard questions.** These are designed to challenge students at the *A/A+* end of the spectrum. Sometimes they require more careful thought than other questions; sometimes they just require independent thought — for example, being able to work through a question from beginning to end, without getting much guidance. Put a lot of effort into practising these if you are aiming for *A* or *A+*.

On the real exam, **there are no difficulty gradings**. My informal percentage gradings are approximately: 15% *E*, 60% *M*, and 25% *H*: but remember that these gradings are geared towards 325 students.

Try to be calm, thorough and careful from the very first question. From past experience, lots of students throw away easy marks at the beginning of the exam as they settle in. If you have spare time, **check your answers, including** any questions you thought were easy. These are the easiest marks to pick up by checking, so they are the most useful checks to make!

Note that easy marks can be interspersed with hard marks, so **always try to attempt all parts of each question**. The marks allotted to each question give a clue about how much work is expected.

Some of the Mock Exams have a substantially higher component of *H* questions: this is to provide extra practice. Your exam is of a similar standard to the exams from 2006 to 2010.

If you find the past exams hard to complete in 3 hours, you should practice several times the 'predictable' types of questions so that you are fully up to speed. Look at the model solutions on the webpage: your answers should contain about the same level of detail. If you put in much more detail, you will waste time. If you put in much less, you will be at risk of losing marks for working. People who skip too many lines of working often make mistakes, and then there is **nothing correct written down** that allows me to award part-marks — make sure you're not in this position!

If you get stuck...

If you get stuck in the middle of a question, **leave that part and keep going**. There is usually enough information given so that you can do the later parts of a question even if you don't get the answers to the earlier parts. If not, then wrong answers will be carried through, so you won't lose marks for a later part by guessing – or just making up – the answer to an earlier part and carrying on. But please be sure to **tell me** what you're doing — a quick written note, like '*Can't do (b), assume answer is 0.2*', is fine.

If you realise something has gone wrong in a question but can't work out where, **let me know**. If you've just made a numerical or transcription error, you won't lose many (if any) marks. But you

are very likely to lose marks if you fiddle your answers, or leave a ludicrous answer (e.g. a negative probability) without a note to say you realise it can't be right. Remember that Stats and Maths are about **communication** as much as thought and reasoning — you need to **communicate** with the person reading your script by using helpful words like “So. . .” or “But. . .”, or “LHS = . . .”, which make all the difference in understanding what you are doing.

Targeted Revision for 2011

The exam revisits several of the questions/ideas covered in the Assignments. Being able to reproduce the assignment questions will not be enough: you need to **understand** how the questions work so that you can apply your understanding in different (but similar) situations on the exam paper. Here is some advice for targeted revision for the 2011 exam. There will be a greater emphasis on named processes on this paper than there has been in the past.

- Assignment 2: Model for gene spread.
- Assignment 3: Two-armed bandit process.
- Assignment 4: Equilibrium in the birth-and-death chain; Extinction-colonisation model.
- Assignment 5: Random Walk with PGFs / defective random variables.
- Mathematical induction: revise the harder examples of induction from the coursebook, Sec 4.3.
- Branching processes: understand that Z_n is a randomly stopped sum, as described on p. 167.

Remember that questions that have been common in the past exams are also likely to occur.

In 2011, there won't be proofs of basic probability manipulations, e.g. there will *not* be a question like Q1 on the 2006 exam. There will *not* be much calculation / computation / solving simultaneous equations: probably less than on previous papers. The exam tests the later chapters most heavily (chapter 4 onwards) because chapters 2 and 3 were already covered in the test.

Revision List

There is not much in the way of factual information that you have to learn. Most of the course is about testing understanding.

Basics

1. Formula for solving a quadratic equation: $ax^2 + bx + c = 0 \Rightarrow x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$.
2. Principle of mathematical induction. Use the Maths Tutor for extra practice: link from the course website.
3. Manipulation of exponents: e.g. $(e^{-x})^2 = e^{-2x}$ while $e^x e^{2x} = e^{x+2x} = e^{3x}$
4. Differentiation using the product rule and the chain rule.
5. Integration of basic functions including x^n , e^{ax} , $\frac{1}{x}$.

Probability

1. Complement: $\mathbb{P}(\bar{A}) = 1 - \mathbb{P}(A)$.
2. Conditional probability: $\mathbb{P}(A | B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}$.
3. Bayes Rule: $\mathbb{P}(A | B) = \frac{\mathbb{P}(B | A)\mathbb{P}(A)}{\mathbb{P}(B)}$.
4. Partition Rule: $\mathbb{P}(B) = \mathbb{P}(B \cap A_1) + \dots + \mathbb{P}(B \cap A_k)$ if A_1, \dots, A_k form a partition of the sample space. Equivalently, $\mathbb{P}(B) = \mathbb{P}(B | A_1)\mathbb{P}(A_1) + \dots + \mathbb{P}(B | A_k)\mathbb{P}(A_k)$.
5. Probability of a union: $\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B)$.
6. Independence: $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$ for independent events A and B .
7. Chains of events: $\mathbb{P}(A_1 \cap A_2 \cap A_3) = \mathbb{P}(A_1)\mathbb{P}(A_2 | A_1)\mathbb{P}(A_3 | A_2 \cap A_1)$.

Discrete distributions

All required information about the Geometric, Binomial, and Poisson distributions is provided in the Attachment. You might need to recognise situations where a Geometric (number of failures before the first success) or Binomial (number of successes in a fixed number of trials) is required. This information is also provided (see Attachment).

Continuous distributions

Where required, the p.d.f., distribution function and mean of the Uniform and Exponential distributions will be provided. Nothing else is required.

Random variables

1. Expectation: $\mathbb{E}(X) = \sum_x x\mathbb{P}(X = x)$ when X is discrete;
 $\mathbb{E}(X) = \int_{-\infty}^{\infty} x f_X(x) dx$ when X is continuous.
2. Expectation of a function of X : $\mathbb{E}(g(X)) = \sum_x g(x)\mathbb{P}(X = x)$ when X is discrete;
 $\mathbb{E}(g(X)) = \int_{-\infty}^{\infty} g(x)f_X(x) dx$ when X is continuous.
3. Variance: $\text{Var}(X) = \mathbb{E}(X^2) - (\mathbb{E}X)^2$ when X is continuous or discrete.
4. Probability as a conditional expectation: for any **discrete** random variable Y ,

$$\mathbb{P}(X \in S) = \sum_y \mathbb{P}(X \in S | Y = y)\mathbb{P}(Y = y).$$

For any **continuous** random variable Y ,

$$\mathbb{P}(X \in S) = \int_y \mathbb{P}(X \in S | Y = y)f_Y(y) dy.$$

5. Independence: for independent X and Y , $\mathbb{E}(XY) = \mathbb{E}(X)\mathbb{E}(Y)$,
 $\mathbb{E}(g(X)h(Y)) = \mathbb{E}(g(X))\mathbb{E}(h(Y))$, and $\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y)$.
6. Conditional expectation: $\mathbb{E}(X | Y = y) = \sum_x x\mathbb{P}(X = x | Y = y)$ when X is discrete.
7. Law of Total Expectation: $\mathbb{E}(X) = \mathbb{E}_Y(\mathbb{E}(X | Y))$.
8. Law of Total Variance: $\text{Var}(X) = \mathbb{E}_Y(\text{Var}(X | Y)) + \text{Var}_Y(\mathbb{E}(X | Y))$.

Probability generating functions

1. Probability generating function: $G_X(s) = \mathbb{E}(s^X)$, if X is discrete and takes values $0, 1, 2, \dots$
2. $P(X = 0) = G_X(0)$ for any random variable X .
3. Defective random variables: X is defective (can take the value ∞) if and only if $G_X(1) < 1$. If X is defective, then $\mathbb{P}(X = \infty) = 1 - G_X(1)$.

Other information is provided on the Attachment.

Branching processes

1. $G_n(s) = G_{n-1}(G(s)) = G(G_{n-1}(s))$, where $G_n(s)$ is the PGF of Z_n , and $G(s)$ is the PGF of family size, Y . This is the *Branching Process Recursion Formula*.
2. Probability of extinction, γ , is the smallest non-negative solution to the equation $G(s) = s$.
3. Conditions for ultimate extinction. Let $\mu = \mathbb{E}(Y)$, then extinction is certain if $\mu \leq 1$. Extinction is less than certain (but generally possible) if $\mu > 1$.

Markov chains

1. Markov property:

$$\mathbb{P}(X_{t+1} = s \mid X_t = s_t, X_{t-1} = s_{t-1}, \dots, X_0 = s_0) = \mathbb{P}(X_{t+1} = s \mid X_t = s_t),$$

for all $t = 1, 2, 3, \dots$ and for all states s_0, s_1, \dots, s_t, s .

2. Distribution of X_t : if $X_0 \sim \pi^T$, then $X_t \sim \pi^T P^t$.
3. Elements of P^t : $\mathbb{P}(X_t = j \mid X_0 = i) = (P^t)_{ij}$.
4. Hitting probabilities: if h_{iA} is the probability of hitting set A , starting from state i , then the vector of hitting probabilities $\mathbf{h}_A = (h_{iA} : i \in S)$ is the minimal non-negative solution to the following first-step analysis equations:

$$h_{iA} = \begin{cases} 1 & \text{for } i \in A, \\ \sum_{j \in S} p_{ij} h_{jA} & \text{for } i \notin A. \end{cases}$$

In practice, it is often easy to see that $h_{kA} = 0$ for some states k by inspection.

5. Mean hitting times: if m_{iA} is the expected number of steps required to hit set A , starting from state i , then the vector of mean hitting times $\mathbf{m}_A = (m_{iA} : i \in S)$ is the minimal non-negative solution to the following first-step analysis equations:

$$m_{iA} = \begin{cases} 0 & \text{for } i \in A, \\ 1 + \sum_{j \notin A} p_{ij} m_{jA} & \text{for } i \notin A. \end{cases}$$

6. Equilibrium distribution: π is an equilibrium distribution for transition matrix P if $\pi^T P = \pi^T$ and $\sum_i \pi_i = 1$.

7. Convergence to equilibrium: if P is irreducible and aperiodic, and an equilibrium distribution exists, then the distribution of X_t converges to the equilibrium distribution as $t \rightarrow \infty$, regardless of its start state.

If an equilibrium distribution exists but the irreducible and aperiodic criteria are **not** satisfied, X_t might or might not converge to equilibrium regardless of start state. In such a case, you will have to work by common sense and inspection.

If X_t does converge, then the long-run probability that X_t is in a particular state is given by the equilibrium distribution. If you are asked for a long-run probability, **you must justify your answer by checking that convergence does occur**. The equilibrium distribution represents long-run probabilities only when the chain **converges** to equilibrium.

Good luck! ☺