MATH 56A SPRING 2008
STOCHASTIC PROCESSES

KIYOSHI IGUSA

These are lecture notes for Math 56a, Spring Semester 2008. Since complete notes from last year are available online, I will write these notes after each class so that they reflect what I actually said or meant to say. These notes are meant replace last year's notes. However, they do not replace the book: "Introduction to Stochastic Processes, 2nd ed." by Greg Lawler. Please read the book with me so we can discuss it. Here is the plan for the course:

(1) st week (jan 17,23) Chap 0: Diff eq's and linear recursion
(2) nd week (jan **24**,28,30,31) Chap 1: Finite Markov chains, Quiz
(3) th week (feb 4,6,7) Chap 1: Finite Markov chains
(4) th week (feb 11,13,14) Chap 2: Countable Markov chains
(5) th week (feb 25,27,28) Chap 3: Continuous time Markov chains, Quiz
(6) th week (mar 3,5,6) Ch 4: Stopping time
(7) th week (mar 10,12,13) Ch 5: Martingales, Quiz
(8) th week (mar 17,19,20) Ch 6: Renewal processes
(9) th week (mar 24,26,27) Ch 7: Reversible Markov chains, Quiz
(10) th week (mar 31, apr 2,3) Ch 8: Weiner process: This is where the intense stuff starts!
(11) th week (apr 7,9,10) Ch 8: more
(12) th week (apr 14,16,17) Ch 9: Stochastic integrals
(13) th week (apr 28,30) Ch 9: more.

The weeks don't quite line up. I need the extra day (Jan 24th) at the end. So, I'll need to start each topic at the end of the previous week.

Date: January 17, 2008.

Introduction

In the first lecture, I discussed the concept of a stochastic process and gave a very quick introduction to two of the main concepts in this course: martingales and Markov processes. I also tried to convey the flavor and philosophy of the course.

A *stochastic process* is defined to be a random process which evolves with time. For example, if you toss two dice then you get the numbers 2 through 12 with a certain fixed probability distribution. This is standard probability theory. An example of a stochastic process might be: Toss two dice and get a total of X_1. Then toss that many dice and get a total of X_2 and so on. As time goes on you will need a lot of dice!

Example. The next example I gave was the question: What is the probability that your family name will survive? The answer I got was 0. I.e., with probability 1, everyone on Earth will have the same last name. This is the male version. You get your last name from your father. But, you get your mitochondria from your mother. The female version is that everybody on Earth will eventually have the same mitochondria which is true!

The setup for the population extinction problem (which we will study more carefully later) is the following.

Start at time $t = 0$ with a male population of N_0.

N_t is the male population after t generations.

X_1 is the number of male offspring from the first man,

X_2 is the number of male offspring from the second man, etc.

Then, the number of males in the next generation will be

$$N_1 = X_1 + X_2 + \cdots + X_{N_0}.$$

Assume that X_i are *independent identically distributed* (i.i.d.) random variables. In particular, they all have the same expected value:

$$\mathbb{E}(X_1) = \mathbb{E}(X_2) = \cdots = \mu.$$

So,

$$\mathbb{E}(N_1) = \mu N_0.$$

This repeats and we get

$$\mathbb{E}(N_t) = \mu^t N_0.$$

This is exponential growth.

Martingale. One thing that is good to do is to make a "martingale":

$$M_t = \frac{N_t}{\mu^t}.$$

Then,

$$\mathbb{E}(M_t) = N_0.$$

This is constant. That makes M_t a martingale. (A *martingale* is a stochastic random variable which you expect to have the same value tomorrow as it has today.)

The "Martingale Convergence Theorem" now tells us that M_t converges to M_∞. On page 119 of our book, it says that, if $\mu > 1$ then $\mathbb{E}(M_\infty) = \mathbb{E}(M_0)$ ($= N_0$ in this case). But, in class I argued that $\mu = 1$ and that the expected value of M_∞ will be zero!

Markov process. For this I converted the problem into a *Markov process.* This is defined to be a system in which there is a fixed set of states and each state there is a fixed probability of going to each other state. For example, in the *random walk* the states are the integer points on the real line. If you are at any point, the probability of going to the left one space is 1/2 and the probability of going to the right one space is 1/2.

For the population problem, the states are: $0, 1, 2, 3, 4, \cdots$ and you are in state N_t at time t. Given N_t men in generation t, there is a certain probability of every possible number of males in the next generation. So, we have a Markov process. We will learn that, in a Markov system, there are only two types of states: "recurrent" and "transient". A *recurrent* state is one that you keep coming back to with probability one. The only recurrent state is 0 (extinction). All other states must be *transient* which means you only go there a finite number of times. I will explain this later in the course.

Since all finite states except 0 are transient, the Markov process will "almost surely" (a.s.) go to 0 or ∞. *Almost surely* means "with probability one." I did not explain why infinity is not possible. In any case, the answer I got was $\mathbb{P}(M_\infty) = 0, a.s.$

Goal of the course. We will look more carefully at this and other example. But the main example that I am interested in is the *Black-Scholes equation.* Some of you already know this equation from economics where it is usually derived using a binomial distribution. We will do a more serious analysis of this equation using *stochastic integration.* Since this is the last topic in the book, we need to cover *the entire book*! We will go very fast, skipping some of the things at the beginning so that we can get to the end.