Math 563 - Fall '21 - Solutions for Final

Do 5 of the 6 problems. Do not turn in 6 problems.

1. Let X_n be an i.i.d. sequence of random variables. Each X_n has an exponential distribution with parameter λ. So its density is $\lambda \exp(-\lambda x)$ for $x \geq 0$. Prove that

$$\limsup_{n\to\infty} \frac{X_n}{\ln n} = \frac{1}{\lambda} \; a.s.$$

Hint: Given a $c > 0$, what can you say about $P(\frac{X_n}{\ln n} \geq c \; i.o.)$.

Solution: Compute

$$P(\frac{X_n}{\ln n} \geq c) = \int_{c \ln n}^{\infty} \lambda e^{-\lambda x} dx = n^{-\lambda c}$$

So $\sum_n P(\frac{X_n}{\ln n} \geq c)$ is ∞ if $\lambda c \leq 1$ and is finite if $\lambda c > 1$. So by the Borel-Cantelli lemma, $P(\frac{X_n}{\ln n} \geq c \; i.o.) = 1$ if $\lambda c \leq 1$ and equals 0 if $\lambda c > 1$. In particular $P(\frac{X_n}{\ln n} \geq 1/\lambda \; i.o.) = 1$. If $\frac{X_n}{\ln n}$ is $\geq 1/\lambda$ infinitely often, then $\limsup_n \frac{X_n}{\ln n}$ is $\geq 1/\lambda$. So $P(\limsup_n \frac{X_n}{\ln n} \geq 1/\lambda) = 1$.

Now let $c > 1/\lambda$. So $P(\frac{X_n}{\ln n} \geq c \; i.o.) = 0$, i.e., with probability one, $\frac{X_n}{\ln n} \geq c$ only a finite number of times. If $\frac{X_n}{\ln n} \geq c$ only a finite number of times, then $\limsup_n \frac{X_n}{\ln n} \leq c$. So $P(\limsup_n \frac{X_n}{\ln n} \leq c) = 1$. This is true for all $c > 1/\lambda$, but note that the exceptional set of probability zero can depend on c. Consider the sequence $c_k = 1/\lambda + 1/k$. The events $E_k = \{\limsup_n \frac{X_n}{\ln n} \leq c_k\}$ are decreasing, and their intersection is the event $\{\limsup_n \frac{X_n}{\ln n} \leq 1/\lambda\}$. By continuity of the probability measure

$$P(\limsup_n \frac{X_n}{\ln n} \leq 1/\lambda) = \lim_k P(\limsup_n \frac{X_n}{\ln n} \leq c_k) = 1$$

Combining this with the result at the end of the previous paragraph yields $P(\limsup_n \frac{X_n}{\ln n} = 1/\lambda) = 1$.

2. Let A_n be a sequence of independent events such that $\lim_{n\to\infty} P(A_n)$ exists and is not 0 or 1. Prove that

$$\frac{1}{\sqrt{n}} \sum_{k=1}^{n} [1_{A_k} - P(A_k)]$$

converges in distribution to a normal random variable and find the mean and variance of this limiting normal random variable.
Solution: Let $p = \lim_n P(A_n)$. We use the Lindeberg-Feller theorem. Let

$$X_{n,m} = \frac{1}{\sqrt{n}}[1_{A_m} - P(A_m)]$$

Note that $E[X_{n,m}] = 0$. Since the A_n are independent, for each n the family $\{X_{n,m} : 1 \le m \le n\}$ is independent. We need to check the conditions for the theorem:

$$EX_{n,m}^2 = \frac{1}{n} P(A_m)[1 - P(A_m)]$$

Since $P(A_m) \to p$, we have

$$\sum_{m=1}^{n} EX_{n,m}^2 = \frac{1}{n} \sum_{m=1}^{n} P(A_m)[1 - P(A_m)] \to p(1-p)$$

Now let $\epsilon > 0$ and consider the indicator function $1_{|X_{n,m}| \ge \epsilon}$.
Since $|1_{A_{n,m}} - P(A_{n,m})| \le 1$, for n sufficiently large (independent of m) this indicator function is zero. So the condition

$$\sum_{m=1}^{n} E[X_{n,m}^2 1_{|X_{n,m}| \ge \epsilon}] \to 0$$

is satisfied. Thus we can apply the LF theorem and conclude

$$\frac{1}{\sqrt{n}} \sum_{k=1}^{n} [1_{A_k} - P(A_k)]$$

converges in distribution to a normal random variable with mean 0 and variance $p(1-p)$.

3. Let X_n and X be random variables on the same probability space. Assume that $E[|X_n|] < \infty$ and $E[|X|] < \infty$.

Here are four types of convergence: (1) pointwise convergence almost surely, (2) convergence in probability, (3) convergence in distribution, and (4) convergence in L^1, i.e., $E|X_n - X|$ converges to zero.

(a) For each of the four types of convergence prove or disprove that if X_n converges to X then $E[X_n]$ converges to $E[X]$.

Solution: None of the first three types of convergence imply $E[X_n]$ converges to $E[X]$. Let $\Omega = [0,1]$ with Lebesgue measure for the probability measure. Let $X_n = n1_{[0,1/n]}$. Then X_n converges to 0 a.s. But $E[X_n] = 1$ while $E[0] = 0$. Since converges almost surely implies the other two forms of convergence this shows that none of the first three types of convegence implies convergence of the expected values. Convergence in L^1 does imply it since

$$|E[X_n] - E[X]| = |E[X_n - X]| \leq E[|X_n - X|]$$

(b) For each of the four types of convergence prove or disprove that if X_n converges to X then $E[\cos(X_n)]$ converges to $E[\cos(X)]$.

Solution: The function $f(x) = \cos(x)$ is bounded and continuous. So convergence in distribution implies $E[\cos(X_n)]$ converges to $E[\cos(X)]$. Since convergence a.s. and convergence in probability both imply convergence in distribution, it follows that any one of the first three types of convergence implies convergence of the expected values. By Chebyshev's inequality for any $\epsilon > 0$,

$$P(|X_n - X| \geq \epsilon) \leq \frac{1}{\epsilon} E[|X_n - X|]$$

So convergence in L^1 implies convergence in probability and so implies $E[\cos(X_n)]$ converges to $E[\cos(X)]$.

4. A random variable X has a geometric distribution if it takes values in the non-negative integers and $P(X = k) = (1-p)^k p, k = 0, 1, 2, \cdots$. Here p is a parameter between 0 and 1. Suppose you have a coin with probability of heads equal to p. Flip it until you get heads for the first time. Let X be the total number of flips. Then X has a geometric distribution.

Let $\alpha > 0$. Let X_n be an independent sequence of random variables such that X_n has the geometric distribution with $p = \alpha/n$. Let $Y_n = X_n/n$. Prove that Y_n converges in distribution and find the limiting distribution.

Solution: Let $\phi_n(t)$ be the characteristic function of Y_n. The characteristic function of the geometric distribution with parameter p is

$$\frac{pe^{it}}{1 - (1-p)e^{it}} = \frac{pe^{it}}{1 - e^{it} + pe^{it}}$$

So

$$\phi_n(t) = \frac{\frac{\alpha}{n} e^{it/n}}{1 - e^{it/n} + \frac{\alpha}{n} e^{it/n}} = \frac{\alpha e^{it/n}}{n(1 - e^{it/n}) + \alpha e^{it/n}}$$

For all t, as $n \to \infty$ this converges to $\frac{\alpha}{\alpha - it}$. This is the characteristic function of an exponential distribution with mean $1/\alpha$.

By the continuity theorem the point-wise convergence of the characteristic functions implies that Y_n converges in distribution to an exponential distribution with mean $1/\alpha$.

5. Let ξ_n be i.i.d. RV's. We assume they are integrable. Let $X_n = \sum_{k=1}^n \xi_k$. Assume that $E[\xi_n] \neq 0$. Recall that this implies X_n is not a martingale. Define $Z_n = \exp(\alpha X_n)$ where α is a constant.
(a) Show that if there is an α such that $\exp(\alpha \xi_n)$ is integrable and

$$E[\exp(\alpha \xi_n)] = 1$$

then Z_n is a martingale.

Solution: We use $Z_{n+1} = Z_n \exp(\alpha \xi_{n+1})$. So

$$E[Z_{n+1}|\mathcal{F}_n] = E[Z_n \exp(\alpha \xi_{n+1})|\mathcal{F}_n]$$

Since Z_n is $\mathcal{F}_n$ measurable, this is

$$= Z_n E[\exp(\alpha \xi_{n+1})|\mathcal{F}_n] = Z_n E[\exp(\alpha \xi_{n+1})]$$

This last inequality follows from the independence of ξ_{n+1} and $\mathcal{F}_n$. This equals Z_n by the hypothesis, so Z_n is a martingale.
(b) Now suppose ξ_n only takes on the values 1 and -1 and let $p = P(\xi_n = 1)$. Now X_n is a non-symmetric random walk. Let a, b be integers with $a < 0 < b$. We start the walk at 0 and X_n is the position at time n. Let τ be the time when we first hit a or b. Assume that the hypotheses of the optional stopping theorem are satisfied. Find the probability the walk hits a before it hits b, i.e., find $P(X_\tau = a)$.

Solution: Solve for α:

$$1 = E[\exp(\alpha \xi_n)] = pe^\alpha + (1-p)e^{-\alpha}$$

So

$$e^\alpha = \frac{1-p}{p}$$

The optional stopping theorem applied to Z_n says $EZ_\tau = EZ_0 = 1$. At the stopping time, Z_n is either a or b. So we have

$$EZ_\tau = e^{\alpha a} P(X_\tau = a) + e^{\alpha b} P(X_\tau = b) = e^{\alpha a} P(X_\tau = a) + e^{\alpha b}(1 - P(X_\tau = a))$$

Solving for $P(X_\tau = a)$ we have

$$P(X_\tau = a) = \frac{1 - e^{\alpha b}}{e^{\alpha a} - e^{\alpha b}} = \frac{1 - \left(\frac{1-p}{p}\right)^b}{\left(\frac{1-p}{p}\right)^a - \left(\frac{1-p}{p}\right)^b}$$

6. Let X be a random variable with $E[X^2] < \infty$. The conditional variance of X given $\mathcal{F}$ is the random variable defined by

$$Var(X|\mathcal{F}) = E[X^2|\mathcal{F}] - (E[X|\mathcal{F}])^2$$

Let X_k be an independent sequence of random variables with $E[X_k^2] < \infty$. Define $S_n = \sum_{k=1}^n X_k$, and let $\mathcal{F}_n$ be the σ-field generated by $X_1, X_2, \cdots, X_n$. Prove that $Var(S_{n+1}|\mathcal{F}_n) = Var(X_{n+1})$ a.s.

Solution: Let μ be the mean of X_{n+1} and σ^2 its variance. Note that $EX_{n+1}^2 = \sigma^2 + \mu^2$.

$$E[S_{n+1}|\mathcal{F}_n] = E[S_n + X_{n+1}|\mathcal{F}_n] = S_n + \mu$$

and

$$E[S_{n+1}^2|\mathcal{F}_n] = E[(S_n + X_{n+1})^2|\mathcal{F}_n] = E[S_n^2 + 2S_n X_{n+1} + X_{n+1}^2|\mathcal{F}_n] = S_n^2 + 2S_n\mu + \sigma^2 + \mu^2$$

So

$$Var(S_{n+1}|\mathcal{F}_n) = S_n^2 + 2S_n\mu + \sigma^2 + \mu^2 - (S_n + \mu)^2 = \sigma^2$$