Homework due: Tuesday 2/7/12

Problems

1. Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space, X be a random variable on the space, and $\mathcal{G}_1 \subseteq \mathcal{G}_2 \subseteq \mathcal{F}$ be two sub-σ-algebras of $\mathcal{F}$. Prove that

$$\mathbf{E}\left(\mathbf{E}\left(X \,|\, \mathcal{G}_2\right) \,|\, \mathcal{G}_1\right) = \mathbf{E}\left(X \,|\, \mathcal{G}_1\right),$$
$$\mathbf{E}\left(\mathbf{E}\left(X \,|\, \mathcal{G}_1\right) \,|\, \mathcal{G}_2\right) = \mathbf{E}\left(X \,|\, \mathcal{G}_1\right).$$

2. Let $S_n = \sum_{k=1}^n X_k$ be the simple symmetric random walk, i.e., the X_k are i.i.d. random variables with $\mathbf{P}(X_k = -1) = \mathbf{P}(X_k = 1) = \frac{1}{2}$. Show that $S_n^2 - n$ is a martingale with respect to the natural filtration $\mathcal{G}_n = \sigma(X_1, \ldots, X_n)$, $n \geq 1$.

3. In the game of **martingale joust**, two teams of knights, the Reds and the Blues, compete to win a large prize. At each turn, the Reds choose one of their knights to joust against one of the Blue knights. Each knight has a positive real number x called his **power** representing his jousting prowess. When two knights A and B with respective powers x and y compete, A wins with probability $x/(x+y)$, and B wins with the complementary probability $y/(x+y)$. The winner gets a boost to their self-confidence and skill level, so that his power becomes $x+y$. The loser is removed from the competition.

   The winning team is the last one to have any active knights remaining. Denote by $x_1, \ldots, x_n$ the powers of the Red knights, and by $y_1, \ldots, y_m$ the powers of the Blue knights at the beginning of the game. Prove that the probability for the Reds to win the competition is exactly

$$\frac{x_1 + \ldots + x_n}{x_1 + \ldots + x_n + y_1 + \ldots + y_m},$$

   regardless of the strategy they employ (the strategy refers to the decision who to send into battle at any round).

4. Let $\Omega = \{1, 2, \ldots, m\}$. Let P and Q be two non-identical probability measures on Ω, each described by a probability vector $p = (p_1, \ldots, p_m)$ and $q = (q_1, \ldots, q_m)$ (i.e., $p_i = P(\{i\})$, $q_i = Q(\{i\})$). Assume that all the p_j's and q_j's are positive. According to the discussion in class, since $P << Q$ and $Q << P$ (the measures are mutually absolutely continuous), given a probabilistic experiment that produces an outcome in Ω and is known to be governed by either the measure P or Q, one can never be sure which of the measures is the correct one. However, in this problem we will show that by performing repeated experiments one can become convinced of the correct answer at an exponentially fast rate.

   (a) Assume that P is the correct measure. Let $X_1, X_2, \ldots$ be an i.i.d. sequence of random samples from P. An experimenter looks at the sample up to time n, and computes the **likelihood ratio**

$$L_n = \frac{P^{\otimes n}(X_1, \ldots, X_n)}{Q^{\otimes n}(X_1, \ldots, X_n)}$$

(i.e., the ratio of how likely it would be to observe that sequence if the measure P governed the process to how likely it would be if the measure Q governed the process). We expect that L_n becomes very large as the size of the sample increases, leading the experimenter to guess with an increasing level of certainty that P is the correct measure.

To make this idea precise, prove that

$$\frac{1}{n}\log L_n \xrightarrow{P\text{-a.s.}} d(P,Q) \text{ as } n \to \infty,$$

where $d(P,Q)$ is the quantity

$$d(P,Q) = \sum_{k=1}^{m} p_k \log\left(\frac{p_k}{q_k}\right).$$

It is not difficult to show that $d(P,Q) > 0$ unless $P \equiv Q$. Thus, over time the experimenter will become increasingly convinced (at an exponential rate, by an average factor of $\exp(d(P,Q))$ per additional sample) that the measure P is the correct measure. The quantity $d(P,Q)$ is known as the **Kullback-Leibler divergence between P and Q**, or the **relative entropy of Q with respect to P**.

(b) To illustrate the idea of (a), assume that $\Omega = \{1,2\}$, $p = \left(\frac{1}{2}, \frac{1}{2}\right)$ and $q = (x, 1-x)$ (this is the problem of trying to decide if a source of random bits you are given access to is unbiased or has bias x). Use a computer or graphing calculator to plot $d(P,Q)$ as a function of x. Note that as x becomes close to $\frac{1}{2}$, the relative entropy $d(P,Q)$, becomes small, which means it becomes increasingly difficult to distinguish between the two measures.

5. **(Optional.)** Let X and Y be random variables defined on the same probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Recall that $\sigma(X) = \{X^{-1}(A) : A \in \mathcal{B}(\mathbb{R})\}$ is the sub-σ-algebra of $\mathcal{F}$ generated by X. Show that $\sigma(X) \subset \sigma(Y)$ (in this case we say X **is measurable with respect to** Y, or X **is Y-measurable**) if and only if there exists a Borel-measurable function $h : \mathbb{R} \to \mathbb{R}$ such that $X = h(Y)$.

Hint. First, prove the "if" part — it's easy. For the "only if" part, the difficulty is in constructing the function h using just the assumption that events of the form $\{a < X < b\}$ are in $\sigma(Y)$. Prove the claim first in the case when X takes only finitely many values $x_1, \ldots, x_n$. Then, for a general X, construct a sequence of discrete approximations to X, i.e., random variables $(X_n)_{n=1}^{\infty}$, each taking only finitely many values, such that

$$X = \lim_{n\to\infty} X_n,$$

and such that each X_n is measurable with respect to $\sigma(X)$ (and therefore also measurable w.r.t. $\sigma(Y)$). Deduce that X_n can be represented as $X_n = h_n(Y)$ for some Borel-measurable function $h_n : \mathbb{R} \to \mathbb{R}$, and use these functions to construct h.