Problem Sheet 6 - Part C Probabilistic Combinatorics - Oxford HT 2008

Chernoff bounds:

1. (i) Let $X_1, X_2, \dots, X_n$ be independent random variables such that $0 \le X_i \le 1$ with proability 1 for all i. Let $S_n = \sum_{i=1}^n X_i$ and let $p = \sum \mathbb{E} X_i/n$, so that $\mathbb{E} S_n = np$. Show that

$$\mathbb{P}(S_n - np \ge t) \le e^{-u(np+t)}(1 - p + pe^u)^n$$

for any $u > 0$, $t > 0$. (Hence the "Chernoff"-type bounds proved in week 6 lectures for the case X_i i.i.d. Bernoulli(p) will also apply in this more general case).

   (ii) Let $a_1, \dots, a_2$ be constants and let $c > 0$. Let $Y_1, \dots, Y_n$ be independent random variables such that $a_i \le Y_i \le a_i + c$ with probability 1, for all i. Give (with brief justification) a version of the Chernoff bound for $\mathbb{P}(S_n - \mathbb{E} S_n \ge t)$, where $S_n = \sum_{i=1}^n Y_i$.

2. A *tournament* on a vertex set V is an orientation of the edges of the complete graph on V. So for each pair $i, j \in V$, $i \ne j$, exactly one of the directed edges $(i \to j)$ and $(j \to i)$ is present. (Think of an all-play-all tournament whose players are the set V; the orientation of the edge between i and j indicates who wins the match between i and j).

   Let σ be a permutation of $\{1, 2, \dots, n\}$. The permutation can be seen as a ranking of the players. We say that an *upset* occurs if the edge $(i \to j)$ is present (i.e. j beats i) and $\sigma(i) < \sigma(j)$ (i.e. i is higher ranked than j).

   Show that there exists a tournament on $\{1, 2, \dots, n\}$ such that the difference between the number of upsets and the number of non-upsets is no greater than $2n^{3/2}\sqrt{\log n}$, for all rankings σ. (So no ranking gives a correct prediction for much more than 50% of the matches).

3. Let $H = (V, E)$ be a hypergraph. Let χ be a two-colouring (red/blue) of its vertices.

   The discrepancy of an edge $e \in E$ under the colouring χ is the absolute difference between the number of blue vertices in the edge and the number of red vertices in the edge.

   The discrepancy of H under χ, denoted $\mathrm{disc}(H, \chi)$ is the maximum over all edges e of the discrepancy of e under χ.

   The discrepancy of H, $\mathrm{disc}(H)$, is then defined as $\min_\chi \mathrm{disc}(H, \chi)$.

   [For example, if H is k-uniform, $\mathrm{disc}(H) < k$ iff H is 2-colourable.]

   (i) Show that if H is k-uniform and has k edges, then $\mathrm{disc}(H) \le 2\sqrt{k \log k}$.

   (ii) Show that if H is k-uniform and each edge intersects at most d other edges, then $\mathrm{disc}(H) \le \sqrt{2k \log(12d)}$.

4. Show that the probability that $G(n, 1/2)$ contains a bipartite subgraph with $n^2/8 + n^{3/2}$ edges is $o(1)$.

Extra question – linearity of expectation:

5. *Buffon's needle.* Suppose a floor is made up of parallel strips of wood, each of width 1. A needle of length $l < 1$ is dropped on the floor. What is the probability $p(l)$ that it crosses the boundary between two strips? (If the picture is not clear, Google for "Buffon's needle"!). The answer leads to an experimental method for estimating π.

   The traditional solution is to average over the position of the end of the needle and over the angle at which the needle lies, using a double integral. This is reasonably straightforward, but an alternative and perhaps prettier approach to the calculation is to use the principle of linearity of expectation (which underlies everything we did with the first moment method).

   Instead of $p(l)$, consider $e(l)$, the *expected number* of times that the needle crosses a line.

   - Show that if $l < 1$ then $p(l) = e(l)$.
   - Show that $e(l) = cl$ for some constant c (for all positive l).

   If we replace a straight needle by some other shape with the same length, the probability of crossing a line will change; however, note that the expectation of the number of line-crossings doesn't! – so to calculate $e(l)$ we may consider "needles" which are not straight lines.

   - Find a "needle" for which the expectation is very easy to calculate. Hence find c.

Course webpage: `http://www.stats.ox.ac.uk/~martin/PC.html`