Problem set 4 - Pair-wise independence

out: 18/11/14
due: 1/12/14

The variational distance between two distributions A, B on Λ is $|A - B| = \frac{1}{2}\sum_{x\in\Lambda}|A(x) - B(x)|$. If A and B are two distributions on Λ_A and Λ_B resp., $A \times B$ denotes the distribution on (a, b) obtained by independently picking a according to A and b according to B. (A, B) denotes some random variable on $\Lambda_A \times \Lambda_B$ with marginal distributions A and B.
For $f : \Lambda \to \Gamma$ and A that is distributed over Λ let $f(A)$ denote the distribution on Γ obtained by picking $a \in A$ and outputting $f(a)$.

1. - Prove $|A - B| = \max_{S\subseteq\Lambda} A(S) - B(S)$.
   - Give a counter example to $|(A, B) - (A', B')| \le |A - A'| + |B - B'|$.
   - Prove that $|A \times B - A' \times B'| \le |A - A'| + |B - B'|$.
   - Assume $|(B|A = a) - C| \le \varepsilon$ for every $a \in \Lambda$. Prove that $|(A, B) - A \times C| \le \varepsilon$.
   - Prove that $|f(A) - f(B)| \le |A - B|$.
   - Extend the previous item to probabilistic functions f.
2. Find an explicit distribution $X_1, ..., X_n$ with support size $Poly(n)$ that is sufficient for derandomizing the maximal-IS algorithm we have seen in class.
3. Prove: If $X = (X_1, ..., X_n)$ is k-wise independent and each X_i is boolean then $Supp(X) \ge B(\frac{k}{2}, n)$, where $B(r, n)$ is the number of words of weight at most r in the n dimensional Boolean cube. Note that X is not necessarily flat (i.e., it might give different weights to elements in its support).
4. You are about to play a game where n coins are laid covered on a table and you uncover and take $\frac{2}{3}n$ coins. You are promised that $k < \frac{n}{3}$ of the coins are pure gold and the rest copper. The catch is that you first have to announce your strategy (be it deterministic or probabilistic) and only then an adversary places the coins on the table. Show that:
   - If you use a deterministic strategy, you can guarantee no gold coin.
   - If you use n random coins you can almost certainly get $\Omega(k)$ gold coins. What is the failure probability?
   - Show that with $O(\log n)$ random coins, you can guarantee $\Omega(k)$ gold coins with probability at least $1 - O(\frac{1}{k})$.
5. Let Y be a distribution over the non-negative integers. Prove that $\frac{(\mathbb{E}[Y])^2}{\mathbb{E}[Y^2]} \le \Pr[Y \ne 0] \le \mathbb{E}[Y]$.
6. Let D be a distribution over Λ. $Col(D) = \Pr_{x_1,x_2\in D}[x_1 = x_2]$. $Supp(D) = \{x|D(x) > 0\}$.
   - Prove that if $Col(D) \le \frac{1}{K}$ then $|Supp(D)| \ge K$.
   - Prove that if $Col(D) \le Col(U_\Lambda)(1 + \varepsilon^2)$ then $|D - U_\Lambda| \le \varepsilon$.