Monte Carlo Methods, Fall 2007, Courant Institute, NYU
Homework 2, Due October 16

1. Verify the recurrence relation for $v_k(x)$ in Section 3.2 of the SDE lecture notes. Use this recurrence relation to verify by induction that $Lv_k = -k\mu v_k$.

2. Read the Monte Carlo lecture notes chapter from Scientific Computing to see how to make error bars for simple samplers.

3. Let $f(x) = C\sqrt{1-x^2}$ be the marginal distribution for the pair (X,Y) uniformly distributed within the unit disk $X^2+Y^2 \leq 1$. Show that sampling X by rejection from the one dimensional uniform distribution gives the same algorithm as sampling the pair (X,Y) within the unit disk by rejection as in homework 1.

4. Make a plot of a Monte Carlo estimate of the exponential moment function $Z(\lambda) = E_f\left[e^{\lambda X}\right]$ for λ in the range $0 \leq \lambda \leq 10$. Use one set of samples of f for the whole curve. Plot the estimate and the error bar for a not very large L (sample size) value and a fairly large value (two plots, one for each L value, each having the estimate together with a one standard deviation upper and lower bound. Explain the high relative accuracy of the estimate for λ near zero (hint: estimate the variance of $e^{\lambda X}$ for small λ).

5. It is important to have some check on the correctness of a Monte Carlo computation, or any other computation for that matter. Find an asymptotic approximation for $Z(\lambda)$ and the variance of $e^{\lambda X}$ that is valid in the limit $\lambda \to \infty$. Are your computational results consistent with these approximations? Why does the Monte Carlo estimate of $Z(\lambda)$ have low relative accuracy for large λ?

6. Write a program that estimates $Z'(\lambda)$ by analytically differentiating the Monte Carlo estimate of $Z(\lambda)$. Why is this better than estimating $Z(\lambda)$ and $Z(\lambda+\Delta\lambda)$ with separate samples of f and using a finite difference? Hint: the worse method has bias (not such a big problem) and much more noise. Use this to write a procedure that solves the equation $E_\lambda[X] = s$ for a given s in the range $0 < s < 1$. Test that your program gives approximately the right answer for s close to zero and s close to one, though the latter may be challenging for the computer.

7. Use all this to create an exponential twist importance sampler that estimates $P(X_1 + \cdots + X_n) > sn$ for $n = 30$ and $s = .5$. Compare your estimate to the analytical approximation using Cramer's theorem (with prefactor). The program should not hard wire n or s so that it can be tested easily with other values. Find a smaller value of n for which direct simulation ($\lambda = 0$) is barely feasible and check that the two estimators (direct and importance sampled) give consistent answers.