Quiz: Variable Variance

CMU 15-462/662

The purpose of this quiz is to help build your familiarity with some elementary quantities and terminology used in Monte Carlo rendering, to get a sense of how to code up basic Monte Carlo integration, to estimate variance (which can be useful for *adaptive* sampling...), and to empirically validate some of the claims made in lecture. Throughout we will consider the simple function

$$f(x) := x^2$$

defined on the interval $0 \leq x \leq 1$. Note that (as with all quizzes) you are free to consult any source you like (textbooks, the internet, your friends, *etc.*) to complement the material covered in class. However, you must write up your final answers by yourself.

1. Compute the expected value $E[Y]$ of the random variable $Y := f(X)$, where X is a random variable uniformly distributed on the interval $0 \leq x \leq 1$. For this question you must compute the *exact* answer by hand, *i.e.*, you should **not** use numerical integration. Show your work.

2. Compute the variance $V[Y]$ of the same random variable. For this question you must again compute the *exact* answer. Show your work.

3. Implement a routine that numerically estimates the expected value $E[Y]$ via Monte Carlo, with the following prototype: `double EY( int N )`. Here `N` is the number of samples; the return value should be the estimate. You can assume that you have a function `unitRand()` that returns independent uniformly distributed random values in the interval $[0,1]$. (An implementation of this method is given on the next page.)

4. Implement a routine that numerically estimates the variance $V(Y)$ via Monte Carlo, with the following prototype: `double VY( int N )`. Here `N` is the number of samples; the return value should be the estimate. You must make no more (and no fewer) than N total calls to `rand()`.

5. Run your two routines `EY(N)` and `VY(N)` for $N = 4096$, and confirm that they produce values that closely match your exact answers from questions 1 and 2.

6. Estimate the *variance of the estimator*[1] `EY(N)` for values of $N = 2^k$, for $k = 1, \ldots, 12$. In other words, as N increases what is the expected deviation of the estimate from the true expected value $E[Y]$ calculated in question 1? For each value of N, confirm that the estimated variance is roughly equal to $\frac{1}{N}V(Y)$, where $V[Y]$ is the exact value of the variance calculated in question 2.

[1] Hint: this question is kind of meta! You are using Monte Carlo to estimate the variance of your Monte Carlo estimator. But that's what you need to do in the "real world" (*e.g.*, for adaptive sampling) since typically you won't know the exact expected value or variance.

Hand In

The quiz you hand in should include:

1. The answers to questions 1 and 2, showing your work.
2. A single printout that includes:
   (a) The routines you implemented for questions 3 and 4.
   (b) The numerical values of `EY(N)` and `VY(N)` computed in question 5.
   (c) For each value of N in question 6, the difference between $\frac{1}{N}V(Y)$ and your numerical estimate of this quantity.

Remember to include your name and AndrewID!

Some skeleton code (in C++):

```
#include <cstdlib>
#include <ctime>
using namespace std;

double unitRand()
{
   const double rRandMax = 1./(double)RAND_MAX;
   return rRandMax * (double) rand();
}

double f( double x )
{
   return x*x;
}

double EY( int N )
{
   // your code here!
   return 0;
}

double VY( int N )
{
   // your code here!
   return 0;
}

int main()
{
   srand(time(NULL));

   // your code here!

   return 0;
}
```