CS 590M **February 6, 2020**
Simulation **Spring Semester 2020**
Peter J. Haas

Assignment #1 Solutions

1. Simulation is used extensively on Wall Street to price exotic options and determine optimal trading rules. The "American" option described in this problem is one of the simplest kinds of options; see the 2003 book by Glasserman for further details. We'll first justify the assertion about the form of $\mathrm{EG}(s,i)$. From basic probability, we know that $X_1+\cdots+X_i$ is distributed as $N(i\mu, i\sigma^2)$. Then

$$\begin{aligned}
\mathrm{EG}(s,i) &= \int_{\log(K/s)}^{\infty} (se^z - K)\frac{1}{\sqrt{2\pi\sigma^2 i}}e^{-(z-i\mu)^2/2i\sigma^2}\,dz \\
&= \int_{-b_i}^{\infty} (se^{i\mu+\sqrt{i}\sigma y} - K)\frac{1}{\sqrt{2\pi}}e^{-y^2/2}\,dy \qquad \left[\text{set } y=(z-i\mu)/\sqrt{i}\sigma\right] \\
&= se^{i\alpha}\int_{-b_i}^{\infty}\frac{1}{\sqrt{2\pi}}e^{-(y-\sqrt{i}\sigma)^2/2}\,dy - K\left(1-\Phi(-b_i)\right) \\
&= se^{i\alpha}\int_{-\sqrt{i}\sigma-b_i}^{\infty}\frac{1}{\sqrt{2\pi}}e^{-x^2/2}\,dx - K\Phi(b_i) \qquad \left[\text{set } x=y-\sqrt{i}\sigma\right] \\
&= se^{i\alpha}\left(1-\Phi(-\sqrt{i}\sigma-b_i)\right) - K\Phi(b_i) \\
&= se^{i\alpha}\Phi(\sqrt{i}\sigma+b_i) - K\Phi(b_i).
\end{aligned}$$

We have used the fact that, by symmetry of the normal distribution, $1-\Phi(x)=\Phi(-x)$.

a. See the class web page for a program that simulates the option strategies. Based on 1000 simulation repetitions, the number of required repetitions was estimated as 5342, so we ran 6000 repetitions to get a 95% confidence interval of 43.61 ± 2.1110 = [41.50, 45.72].

b. Based on 1000 simulation repetitions, the number of required repetitions was estimated as 2504, so we ran 3000 repetitions to get a 95% confidence interval of 18.09 ± 0.7990 = [17.29, 18.89].

c. Using 20,000 repetitions, we obtained independent confidence intervals of 43.33 ± 1.1591 for EG and 17.97 ± 0.3064 for TMR. Let $\bar{U}$ and $\bar{V}$ be the average of the gains for EG and TMR, respectively. By the central limit theorem, the two estimators $\bar{U}$ and $\bar{V}$ are statistically independent and approximately normally distributed with means given by 43.31 and 18.00, respectively. The variances are given by $\sigma_{\bar{U}}^2 = H_{\bar{U}}^2 / z_{0.95}^2$ and $\sigma_{\bar{V}}^2 = H_{\bar{V}}^2 / z_{0.95}^2$, where H_X denotes the confidence interval half-width for estimator X. By the hint, $\bar{V}-\bar{U}$ has approximately a normal distribution with mean $43.33-17.97=25.36$ (which is the combined point estimate) and a variance of $\sigma_{\bar{V}-\bar{U}}^2 = \sigma_{\bar{U}}^2 + \sigma_{\bar{V}}^2 = (H_{\bar{U}}^2 + H_{\bar{V}}^2)/z_{0.95}^2$. Thus the combined 95% confidence interval half-width is $z_{0.95}\sigma_{\bar{V}-\bar{U}}$ (by a derivation as in class), so that

$$\mathrm{HW} = z_{0.95}\sigma_{\bar{V}-\bar{U}} = (H_{\bar{U}}^2 + H_{\bar{V}}^2)^{1/2} = (1.1590^2 + 0.3064^2)^{1/2} = 1.1988.$$

d. When using the same set of stock prices to compute the two rewards during each of 20,000 simulation repetitions, we got a 95% confidence interval for $E[U-V]=E[U]-E[V]$ of 24.92 ± 1.0956. The half-width of 1.0956 is about 8.6% shorter than the half-width 0f 1.1988 obtained in part (c). Intuitively, we get a sharper picture of the difference between the two option policies by subjecting each to the exact same sequence of stock fluctuations, so that, in each repetition, any differences are due to the policy behaviors only, and not to luck of the draw.

2. This problem illustrates some basic techniques of Monte Carlo integration.
   a. With $Z_i = h(U_i)$ for $i \ge 1$, we have $E[Z_i] = E[h(U_i)] = \int_{-\infty}^{\infty} h(x) f_U(x)\,dx = \int_0^1 h(x)\,dx = I$ for each i. So the strong law of large numbers implies that $(1/n)\sum_{i=1}^n h(U_i) = (1/n)\sum_{i=1}^n Z_i \to I$ with probability 1 as $n \to \infty$. Thus the proposed algorithm gives a strongly consistent estimate of I. The techniques discussed in class can be used to obtain a confidence interval for I and to choose a suitable value of n to achieve a desired level of absolute or relative precision at a specified level of confidence.
   b. After we make the suggested transformation (see p. 12 of the probability/stats refresher handout for a quick review of how to make a simple change of variable in an integral), we find that $I = \int_0^1 h(a+(b-a)y)(b-a)\,dy = \int_0^1 g(y)\,dy$, where $g(y) = (b-a)h(a+(b-a)y)$. We now proceed as in Part (a), but with h replaced by g.
   c. After making the transformation, we have $I = \int_0^1 g(y)\,dy$, where $g(y) = h(y^{-1}-1)/y^2$. Now proceed as in Part (a).
   d. Take Z as $e^{(U+V)^2}$, where U and V are independent uniform(0,1) random variables. Then $E[Z] = E[e^{(U+V)^2}] = \int_0^1\int_0^1 e^{(x+y)^2}\,dx\,dy$. So proceeding analogously to Part (a), generate $U_1, U_2, \ldots, U_n$ and $V_1, V_2, \ldots, V_n$, and compute $\hat{I}_n = (1/n)\sum_{i=1}^n e^{(U_i+V_i)^2}$ as an estimator of I.
   e. Following the hint, write $I = \int_0^\infty\int_0^\infty g(x,y)\,e^{-(x+y)^2}\,dx\,dy$. Making the transformations $u = 1/(x+1)$ and $v = 1/(y+1)$ as in Part (c), we have
$$\begin{aligned} I &= \int_0^1\int_0^1 u^{-2}v^{-2} g(u^{-1}-1, v^{-1}-1)\, e^{-(u^{-1}+v^{-1}-2)^2}\,du\,dv \\ &= \int_0^1\int_0^1 u^{-2}v^{-2} g(v,u)\, e^{-(u^{-1}+v^{-1}-2)^2}\,du\,dv \end{aligned}$$
   So generate $U_1, U_2, \ldots, U_n$ and $V_1, V_2, \ldots, V_n$, and set $\hat{I}_n = (1/n)\sum_{i=1}^n U_i^{-2}\,V_i^{-2} g(V_i, U_i)\, e^{-(U_i^{-1}+V_i^{-1}-2)^2}$.

3. This result will also be needed later in the course, when we study random-number generation. We have $f_X(x) = I(0 \le x \le 1)$, $f_Y(y) = I(0 \le y \le 1)$, and
$$F_X(x) = \begin{cases} 0 & \text{if } x \le 0 \\ x & \text{if } 0 \le x \le 1 \\ 1 & \text{if } x \ge 1 \end{cases}$$
By the hint, we have $F_Z(z) = \int_0^1 F_X(z-y)\,dy$. (The last step in the hint uses the independence of X and Y.) If $z \in [0,1]$, then
$$F_Z(z) = \int_0^z (z-y)\,dy = \left.\frac{-(z-y)^2}{2}\right|_0^z = z^2/2.$$
If $z \in [1,2]$, then
$$F_Z(z) = \int_0^{z-1} 1\,dy + \int_{z-1}^1 (z-y)\,dy = z-1+\left.\frac{-(z-y)^2}{2}\right|_{z-1}^1 = 2z - \frac{z^2}{2} - 1.$$ Taking derivatives,

$$f_Z(z) = \begin{cases} z & \text{if } 0 \le z \le 1 \\ 2-z & \text{if } 1 \le z \le 2 \\ 0 & \text{otherwise} \end{cases}$$

The density function for the (obviously-named) "triangular distribution" is sketched below, along with the histogram produced by the given R code.

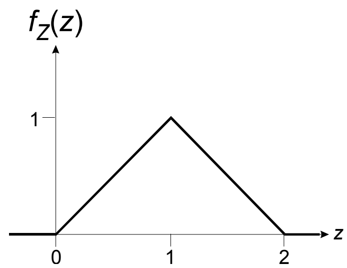

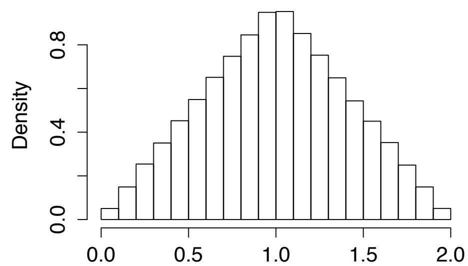

4. The estimation procedure introduced in this problem is known as Latin hypercube sampling (LHS), because it is based on "Latin squares."
   a. Using R, the standard method produced an estimate of 0.858 with a standard error of 0.007, whereas LHS produced an estimate of 0.851with a standard error of 0.003. (For this simple problem, the exact answer can be numerically computed as 0.8511.) Clearly LHS is the superior method, with a standard error less than half of that for the standard method (and so LHS would yield a confidence interval that is half as wide.)
   b. A typical realization of $\mathbf{W}_1, \mathbf{W}_2, \ldots, \mathbf{W}_5$ is indicated in the figure below.

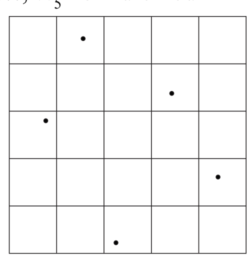

The points form a *Latin square* in that there is exactly one point in each row and each column. Thus the points are systematically spread out, in contrast to ordinary uniform random variables, which can be "clumpy" in that some rows or columns could contain multiple points and others could contain none. Repeating this process *n* times yields an overall set of points such that there are exactly *n* points in each row and column. This relatively uniform distribution of points covers the domain of the function *h* more thoroughly, yielding better Monte Carlo estimates. Note that we could achieve even more thorough coverage by a "complete stratification" scheme in which we place N / K^2 points uniformly in each of the K^2 subregions. For *d*-dimensional problems with $d \gg 2$, as is typical in situations where we would actually want to use Monte Carlo integration, complete stratification would require at least K^d points, which could be a huge number. LHS is a good compromise strategy, in which a weaker, but still effective, type of uniformity is achieved at a much lower cost. "Quasi-random number" techniques produce points that are even more uniformly distributed, but these techniques are quite complicated and beyond the scope of the course.

c. Using the simple formula would only be correct if, for each *i*, the random variables $\{h(\mathbf{W}_k^{(i)}):\ k=1,2,\ldots,K\}$ that make up Q_i were mutually independent. However, each set $\mathbf{W}_1^{(i)},\mathbf{W}_2^{(i)},\ldots,\mathbf{W}_K^{(i)}$ of random variables is highly dependent, since they must lie in disjoint rows and columns. The random variables $Q_1,Q_2,\ldots,Q_n$ are i.i.d., however, so the formula given in the presented algorithm is correct.