Monte Carlo analysis

S.M. Iacus

Department of Economics, Management and Quantitative Methods, University of Milan
CREST JST

June 25, 2019

YSS2019

Monte Carlo Method

The Monte Carlo Method is a numeric method based on statistical arguments which can be used to evaluate integrals. In particular, the expected value of a random variable X with density $f(\cdot)$ is an integral of this form

$$\mathbf{E}X = \int x f(x) \mathrm{d}x$$

In finance, we use the Monte Carlo Method to evaluate the expected payoff of some derivative in order to price it.

Monte Carlo Method

Suppose we are interested in calculating the expected value $\mathbf{E}g(Y)$, with g is any function and Y is a given random variable.

Assume we know how to simulate n pseudo-random numbers $y_1, \ldots, y_n$ according to the distribution of Y.

Monte Carlo Method

Suppose we are interested in calculating the expected value $\mathbf{E}g(Y)$, with g is any function and Y is a given random variable.

Assume we know how to simulate n pseudo-random numbers $y_1, \ldots, y_n$ according to the distribution of Y.

Then, we can think to approximate $\mathbf{E}g(Y)$ with the arithmetic mean of the numbers $g(y_i)$

$$\mathbf{E}g(Y) \simeq \frac{1}{n}\sum_{i=1}^{n} g(y_i)$$

Monte Carlo Method

Is this a good approximation?

$$\mathbf{E}g(Y) \simeq \frac{1}{n}\sum_{i=1}^{n} g(y_i) = \bar{g}_n$$

Yes! This is guaranteed by the Law of Large Numbers.

Monte Carlo Method

Is this a good approximation?

$$\mathbf{E}g(Y) \simeq \frac{1}{n}\sum_{i=1}^{n} g(y_i) = \bar{g}_n$$

Yes! This is guaranteed by the Law of Large Numbers.

In addition, by the Central Limit Theorem we also have

$$\frac{1}{n}\sum_{i=1}^{n} g(y_i) \sim N\left(\mathbf{E}g(Y), \frac{1}{n}\mathrm{Var}(g(Y))\right)$$

Properties of the MC method

In summary, the number we estimate by n simulations, will have a precision, around the true value $\mathbf{E}g(Y)$ of order $\epsilon_n = 1/\sqrt{n}$.

Moreover, the estimate $\frac{1}{n}\sum_{i=1}^n g(y_i)$ is contained in the interval

$$(\mathbf{E}g(Y) - \epsilon_n, \mathbf{E}g(Y) + \epsilon_n)$$

with probability 68.3%.

Properties of the MC method

In summary, the number we estimate by n simulations, will have a precision, around the true value $\mathbf{E}g(Y)$ of order $\epsilon_n = 1/\sqrt{n}$.

Moreover, the estimate $\frac{1}{n}\sum_{i=1}^{n} g(y_i)$ is contained in the interval

$$(\mathbf{E}g(Y) - \epsilon_n, \mathbf{E}g(Y) + \epsilon_n)$$

with probability 68.3%. More precisely, a 95% confidence interval will be of this form

$$\left(\mathbf{E}g(X) - 1.96\frac{\sigma}{\sqrt{n}}, \mathbf{E}g(X) + 1.96\frac{\sigma}{\sqrt{n}}\right),$$

with $\sigma = \sqrt{\mathrm{Var} g(X)}$.

Properties of the MC method

The confidence interval depends on $\mathrm{Var} g(X)$ which is estimated as

$$\hat{\sigma}^2 = \frac{1}{n-1}\sum_{i=1}^{n}(g(x_i) - \bar{g}_n)^2$$

and the approximate 95% Monte Carlo confidence interval for $\mathbf{E}g(X)$ is

$$\left(\bar{g}_n - 1.96\frac{\hat{\sigma}}{\sqrt{n}}, \bar{g}_n + 1.96\frac{\hat{\sigma}}{\sqrt{n}}\right).$$

The quantity $\hat{\sigma}/\sqrt{n}$ is called the *standard error* but standard error is itself a random quantity and thus subject to variability; hence one should interpret this value as a “qualitative” measure of accuracy.

An application to integral evaluation

Suppose we have a function g defined on $[a, b]$ with values in $[c, d]$, $c, d \geq 0$. We are interested in

$$\int_a^b g(x)\mathrm{d}x$$

We can make use of the Monte Carlo method in the following way:

An application to integral evaluation

Suppose we have a function g defined on $[a, b]$ with values in $[c, d]$, $c, d \geq 0$. We are interested in

$$\int_a^b g(x)\mathrm{d}x$$

We can make use of the Monte Carlo method in the following way:

- let $s = 0$, $A = (b - a) * (d - c)$

An application to integral evaluation

Suppose we have a function g defined on $[a, b]$ with values in $[c, d]$, $c, d \geq 0$. We are interested in

$$\int_a^b g(x)\mathrm{d}x$$

We can make use of the Monte Carlo method in the following way:

- let $s = 0$, $A = (b - a) * (d - c)$
- generate a pseudo-random number $U_i \sim U(a, b)$

An application to integral evaluation

Suppose we have a function g defined on $[a, b]$ with values in $[c, d]$, $c, d \geq 0$. We are interested in

$$\int_a^b g(x)\mathrm{d}x$$

We can make use of the Monte Carlo method in the following way:

- let $s = 0$, $A = (b - a) * (d - c)$
- generate a pseudo-random number $U_i \sim U(a, b)$
- generate a second pseudo-random number $V_i \sim U(c, d)$

An application to integral evaluation

Suppose we have a function g defined on $[a, b]$ with values in $[c, d]$, $c, d \geq 0$. We are interested in

$$\int_a^b g(x)\mathrm{d}x$$

We can make use of the Monte Carlo method in the following way:

- let $s = 0$, $A = (b - a) * (d - c)$
- generate a pseudo-random number $U_i \sim U(a, b)$
- generate a second pseudo-random number $V_i \sim U(c, d)$
- if $V_i < g(U_i)$ then set $s = s + 1$

An application to integral evaluation

Suppose we have a function g defined on $[a, b]$ with values in $[c, d]$, $c, d \geq 0$. We are interested in

$$\int_a^b g(x)\mathrm{d}x$$

We can make use of the Monte Carlo method in the following way:

- let $s = 0$, $A = (b - a) * (d - c)$
- generate a pseudo-random number $U_i \sim U(a, b)$
- generate a second pseudo-random number $V_i \sim U(c, d)$
- if $V_i < g(U_i)$ then set $s = s + 1$
- iterate n times

An application to integral evaluation

Suppose we have a function g defined on $[a, b]$ with values in $[c, d]$, $c, d \geq 0$. We are interested in

$$\int_a^b g(x)\mathrm{d}x$$

We can make use of the Monte Carlo method in the following way:

- let $s = 0$, $A = (b - a) * (d - c)$
- generate a pseudo-random number $U_i \sim U(a, b)$
- generate a second pseudo-random number $V_i \sim U(c, d)$
- if $V_i < g(U_i)$ then set $s = s + 1$
- iterate n times

The value $\int_a^b g(x)\mathrm{d}x \simeq A * s/n$

Example: $\int_0^2 x^2 \mathrm{d}x = \frac{2^3}{3} = 2.6\bar{6}$

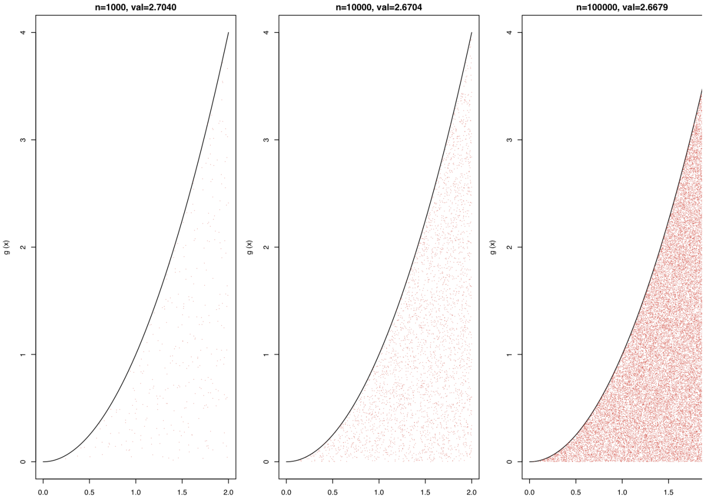

An application to integral evaluation

The algorithm just counts the number of points of random coordinates (u, v) which fall below the curve g.

The proportion of this points which are below the graph of g is s/n.

But the points of coordinates (u, v) have been randomly extracted from the rectangle of area A, thus, the area under the curve g (i.e. the integral of g) is obtained using the formula $A * s/n$.

An application to integral evaluation

```
> set.seed(123)
> g <- function(x) x^2
> a <- 0
> b <- 2
> c <- 0
> d <- 4
> A <- (b - a) * (d - c)
> n <- 1e+05
> x <- runif(n, a, b)
> y <- runif(n, c, d)

> A * sum(y < g(x))/n
[1] 2.66792   # Monte Carlo estimate

> integrate(g, a, b)   # Numerical integration
2.666667 with absolute error < 3.0e-14
```

Connection with the Monte Carlo method

We set $s = s + 1$ whenever the $V_i < g(U_i)$ occurs, with $U_i \sim U(a, b)$ and $V_i \sim U(c, d)$, $i = 1, \dots, n$.

If we set $Y_i = \mathbf{1}_{\{V_i < g(U_i)\}}$, $i = 1, \dots, n$, we obtain a sample of Bernoulli random numbers which are all i.i.d as $Y \sim Ber(p)$, with $p = Pr(\text{“being below the curve } g''\text{)}$.

Then,

$$\bar{Y}_n = \frac{1}{n} \sum_{i=1}^{n} Y_i \simeq \mathbf{E} Y = p$$

and thus

$$A \cdot \bar{Y}_n = A \cdot \frac{1}{n} \sum_{i=1}^{n} Y_i = A \cdot \frac{s}{n} \simeq \int_a^b g(x)\mathrm{d}x$$

Accuracy and speed of convergence

In some cases, Monte Carlo intervals are not very informative if the variance of $Y = g(X)$ is too large.

Accuracy and speed of convergence

In some cases, Monte Carlo intervals are not very informative if the variance of $Y = g(X)$ is too large.

Let $Y = g(X) = e^{\beta X}$ with $X \sim N(0,1)$, and assume we are interested in $\mathbf{E}g(X)$ with $\beta = 5$.

Accuracy and speed of convergence

In some cases, Monte Carlo intervals are not very informative if the variance of $Y = g(X)$ is too large.

Let $Y = g(X) = e^{\beta X}$ with $X \sim N(0,1)$, and assume we are interested in $\mathbf{E}g(X)$ with $\beta = 5$.

The analytical value can be calculated as $e^{\beta^2/2} = 268337.3$, and the true standard deviation $\sigma = \sqrt{e^{2\beta^2} - e^{\beta^2}} = 72004899337$, quite a big number with respect to the mean of Y.

Accuracy and speed of convergence

In some cases, Monte Carlo intervals are not very informative if the variance of $Y = g(X)$ is too large.

Let $Y = g(X) = e^{\beta X}$ with $X \sim N(0,1)$, and assume we are interested in $\mathbf{E}g(X)$ with $\beta = 5$.

The analytical value can be calculated as $e^{\beta^2/2} = 268337.3$, and the true standard deviation $\sigma = \sqrt{e^{2\beta^2} - e^{\beta^2}} = 72004899337$, quite a big number with respect to the mean of Y.

Suppose we want to estimate $\mathbf{E}Y$ via the Monte Carlo method using 100000 replications and construct 95% confidence intervals using the true standard deviation σ and the estimated standard error. Next R code does the job.

Accuracy and speed of convergence

```
> set.seed(123)
> n <- 1000000
> beta <-5
> x <- rnorm(n)
> y <- exp(beta*x)
>
> # true value of E(Y)
> exp(beta^2/2)
[1] 268337.3
> # MC estimation of E(Y)
> mc.mean <- mean(y)
> mc.mean
[1] 199659.2
> mc.sd <- sd(y)
> true.sd <- sqrt(exp(2*beta^2) - exp(beta^2))
>
> # MC conf. interval based on true sigma
> mc.mean - true.sd*1.96/sqrt(n)
[1] -140929943
> mc.mean + true.sd*1.96/sqrt(n)
[1] 141329262
>
> # MC conf. interval based on estimated sigma
> mc.mean - mc.sd*1.96/sqrt(n)
[1] 94515.51
> mc.mean + mc.sd*1.96/sqrt(n)
[1] 304802.9
```

Accuracy and speed of convergence

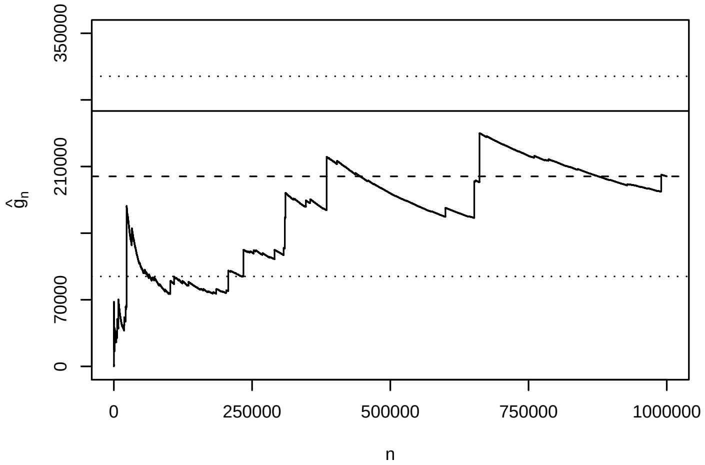

Accuracy and speed of convergence

Running the previous R code we obtain the two intervals

$$(-140929943; 141329262) \quad \text{using } \sigma$$

and

$$(94515.51; 304802.9) \quad \text{using } \hat{\sigma}$$

with an estimated value of $\mathbf{E}g(X)$, $\hat{g}_n = 199659.2$.

As one can see, the confidence interval based on σ contains the true value of $\mathbf{E}g(X)$ but is too large and hence meaningless.

The confidence interval based on $\hat{\sigma}$ is smaller but still large.

The first effect is due to the big variance of $g(X)$, while the second is due to the fact that the sample variance underestimates the true one ($\hat{\sigma} = 53644741$).

Accuracy and speed of convergence

We have seen that variability affects stability of the Monte Carlo method

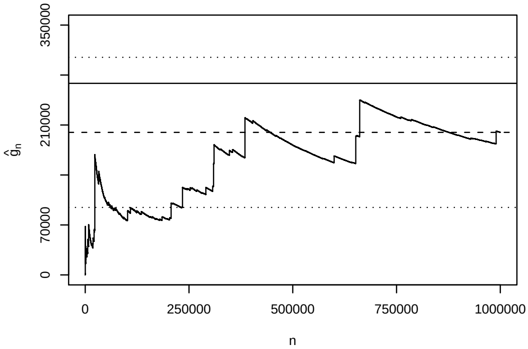

We now present a couple of techniques to constrain instability

Preferential sampling

The idea of this method is to express $\mathbf{E}g(X)$ in a different form in order to reduce its variance.

Preferential sampling

The idea of this method is to express $\mathbf{E}g(X)$ in a different form in order to reduce its variance.

Let $f(\cdot)$ be the density of X; thus

$$\mathbf{E}g(X) = \int_{\mathbb{R}} g(x)f(x)\mathrm{d}x.$$

Introduce now another strictly positive density $h(\cdot)$. Then,

$$\mathbf{E}g(X) = \int_{\mathbb{R}} \frac{g(x)f(x)}{h(x)} h(x)\mathrm{d}x$$

and

$$\mathbf{E}g(X) = \mathbf{E}\left(\frac{g(Y)f(Y)}{h(Y)}\right) = \mathbf{E}\tilde{g}(Y),$$

with Y a random variable with density $h(\cdot)$ and $\tilde{g}(\cdot) = g(\cdot)f(\cdot)/h(\cdot)$

Preferential sampling

Thus

$$\mathbf{E}g(X) = \mathbf{E}\left(\frac{g(Y)f(Y)}{h(Y)}\right) = \mathbf{E}\tilde{g}(Y)$$

If we are able to determine an $h(\cdot)$ such that $\mathrm{Var}\tilde{g}(Y) < \mathrm{Var}g(X)$, then we have reached our goal. But

Preferential sampling

Thus

$$\mathbf{E}g(X) = \mathbf{E}\left(\frac{g(Y)f(Y)}{h(Y)}\right) = \mathbf{E}\tilde{g}(Y)$$

If we are able to determine an $h(\cdot)$ such that $\mathrm{Var}\tilde{g}(Y) < \mathrm{Var}g(X)$, then we have reached our goal. But

$$\mathrm{Var}\tilde{g}(Y) = \mathbf{E}\tilde{g}(Y)^2 - (\mathbf{E}\tilde{g}(Y))^2 = \int_{\mathbb{R}} \frac{g^2(x)f^2(x)}{h(x)}\mathrm{d}x - (\mathbf{E}g(X))^2 .$$

Let $g()$ be strictly positive and choose $h(x) = g(x)f(x)/\mathbf{E}g(X)$, then

$$\mathrm{Var}\tilde{g}(Y) = \mathbf{E}g(X)\int_{\mathbb{R}} \frac{g^2(x)f^2(x)}{g(x)f(x)}\mathrm{d}x - (\mathbf{E}g(X))^2 = 0$$

Preferential sampling

Thus

$$\mathbf{E}g(X) = \mathbf{E}\left(\frac{g(Y)f(Y)}{h(Y)}\right) = \mathbf{E}\tilde{g}(Y)$$

If we are able to determine an $h(\cdot)$ such that $\mathrm{Var}\tilde{g}(Y) < \mathrm{Var}g(X)$, then we have reached our goal. But

$$\mathrm{Var}\tilde{g}(Y) = \mathbf{E}\tilde{g}(Y)^2 - (\mathbf{E}\tilde{g}(Y))^2 = \int_{\mathbb{R}} \frac{g^2(x)f^2(x)}{h(x)}\mathrm{d}x - (\mathbf{E}g(X))^2 .$$

Let $g()$ be strictly positive and choose $h(x) = g(x)f(x)/\mathbf{E}g(X)$, then

$$\mathrm{Var}\tilde{g}(Y) = \mathbf{E}g(X)\int_{\mathbb{R}} \frac{g^2(x)f^2(x)}{g(x)f(x)}\mathrm{d}x - (\mathbf{E}g(X))^2 = 0$$

This is nice only in theory because, of course, we don't know $\mathbf{E}g(X)$.

Preferential sampling

But the expression of $h(x)$ suggests a way to obtain a useful approximation: just take $\tilde{h}(x) = |g(x)f(x)|$ (or something close to it), then normalize it by the value of its integral, and use

$$h(x) = \frac{\tilde{h}(x)}{\int_{\mathbb{R}} \tilde{h}(x)\mathrm{d}x}\,.$$

Of course this is simple to say and hard to solve in specific problems, as integration should be done analytically and not using the Monte Carlo technique again.

Let us see an example.

Preferential sampling

Suppose we want to calculate $\mathbf{E}g(X)$ with

$$g(x) = \max(0, K - e^{\beta x}) = (K - e^{\beta x})_+$$

K and β constants, and $X \sim N(0,1)$.

This is the price of a *put option* in the Black and Scholes framework with explicit solution

$$\mathbf{E}\left(K - e^{\beta X}\right)_+ = K\Phi\left(\frac{\log(K)}{\beta}\right) - e^{\frac{1}{2}\beta^2}\Phi\left(\frac{\log(K)}{\beta} - \beta\right),$$

where Φ is the cumulative distribution function of the standard Gaussian law.

The true value, in the case $K = \beta = 1$, is $\mathbf{E}g(X) = 0.2384217$.

Preferential sampling

Suppose we want to calculate $\mathbf{E}g(X)$ with

$$g(x) = \max(0, K - e^{\beta x}) = (K - e^{\beta x})_+$$

K and β constants, and $X \sim N(0,1)$.

This is the price of a *put option* in the Black and Scholes framework with explicit solution

$$\mathbf{E}\left(K - e^{\beta X}\right)_+ = K\Phi\left(\frac{\log(K)}{\beta}\right) - e^{\frac{1}{2}\beta^2}\Phi\left(\frac{\log(K)}{\beta} - \beta\right),$$

where Φ is the cumulative distribution function of the standard Gaussian law.

The true value, in the case $K = \beta = 1$, is $\mathbf{E}g(X) = 0.2384217$.

Monte Carlo: true value 0.2384217

```
>  set.seed(123)
>  n <- 1000
>  beta <-1
>  K <- 1
>  x <- rnorm(n)
>  y <- sapply(x, function(x) max(0,K-exp(beta*x)))
>
>  # the true value
>  K*pnorm(log(K)/beta)-exp(beta^2/2)*pnorm(log(K)/beta-beta)
[1] 0.2384217
>
> # MC value
> mean(y)
[1] 0.2313868
```

n	$\hat{g}_n$	95% conf. interval
100	0.206	(0.153 ; 0.259)
1000	0.231	(0.213 ; 0.250)
10000	0.238	(0.232 ; 0.244)

Variance reduction

We now try to rewrite $\mathbf{E}g(X)$ as $\mathbf{E}\tilde{g}(Y)$ in order to reduce its variance. Indeed, $\mathbf{E}g(X)$ can be rewritten as

$$\mathbf{E}\left(K - e^{\beta X}\right)_+ = \int_{\mathbb{R}} \frac{(K - e^{\beta x})_+}{\beta|x|}\beta|x|\frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}}\mathrm{d}x,$$

set $K = 1$ and notice that $e^x - 1 \simeq x$ for x close to 0.

Variance reduction

We now try to rewrite $\mathbf{E}g(X)$ as $\mathbf{E}\tilde{g}(Y)$ in order to reduce its variance. Indeed, $\mathbf{E}g(X)$ can be rewritten as

$$\mathbf{E}\left(K - e^{\beta X}\right)_+ = \int_{\mathbb{R}} \frac{(K - e^{\beta x})_+}{\beta|x|}\beta|x|\frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}}\mathrm{d}x,$$

set $K = 1$ and notice that $e^x - 1 \simeq x$ for x close to 0.

By the change of variable $x = \sqrt{y}$ for $x > 0$ and $x = -\sqrt{y}$ for $x < 0$, the integral above can be rewritten as

$$\int_0^\infty \frac{\left(1 - e^{\beta\sqrt{y}}\right)_+ + \left(1 - e^{-\beta\sqrt{y}}\right)_+}{\sqrt{2\pi}\sqrt{y}}\frac{e^{-\frac{1}{2}y}}{2}\mathrm{d}y,$$

Variance reduction

We now try to rewrite $\mathbf{E}g(X)$ as $\mathbf{E}\tilde{g}(Y)$ in order to reduce its variance. Indeed, $\mathbf{E}g(X)$ can be rewritten as

$$\mathbf{E}\left(K - e^{\beta X}\right)_+ = \int_{\mathbb{R}} \frac{(K - e^{\beta x})_+}{\beta|x|} \beta|x| \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \mathrm{d}x,$$

set $K = 1$ and notice that $e^x - 1 \simeq x$ for x close to 0.

By the change of variable $x = \sqrt{y}$ for $x > 0$ and $x = -\sqrt{y}$ for $x < 0$, the integral above can be rewritten as

$$\int_0^\infty \frac{\left(1 - e^{\beta\sqrt{y}}\right)_+ + \left(1 - e^{-\beta\sqrt{y}}\right)_+}{\sqrt{2\pi}\sqrt{y}} \frac{e^{-\frac{1}{2}y}}{2} \mathrm{d}y,$$

i.e.

$$\tilde{g}(x) = \frac{\left(1 - e^{\beta\sqrt{x}}\right)_+ + \left(1 - e^{-\beta\sqrt{x}}\right)_+}{\sqrt{2\pi}\sqrt{x}}$$

Variance reduction

$$\mathbf{E}g(X) = \int_0^\infty \frac{\left(1 - e^{\beta\sqrt{y}}\right)_+ + \left(1 - e^{-\beta\sqrt{y}}\right)_+}{\sqrt{2\pi}\sqrt{y}} \frac{e^{-\frac{1}{2}y}}{2} \mathrm{d}y = \mathbf{E}\tilde{g}(Y)$$

Remark that $f(y) = \lambda e^{-\lambda y}$, with $\lambda = \frac{1}{2}$, is the density of the exponential distribution. Therefore,

$$\mathbf{E}g(X) = \mathbf{E}\left(\frac{\left(1 - e^{\beta\sqrt{Y}}\right)_+ + \left(1 - e^{-\beta\sqrt{Y}}\right)_+}{\sqrt{2\pi}\sqrt{Y}} \right)$$

can be evaluated as the expected value of a function of the exponential random variable Y.

At this point you should believe that this second expression with $\tilde{g}(\cdot)$ has lower variance than the original with $g(\cdot)$.

Monte Carlo with variance reduction: true value 0.2384217

```
>  set.seed(123)
>  n <- 1000
>  beta <-1
>  K <- 1
>
>  x <- rexp(n,rate=0.5)
>  h <- function(x) (max(0,1-exp(beta*sqrt(x))) +
+    max(0,1-exp(-beta*sqrt(x))))/sqrt(2*pi*x)
>  y <- sapply(x, h)
>  # the true value
>  K*pnorm(log(K)/beta)-exp(beta^2/2)*pnorm(log(K)/beta-beta)
[1] 0.2384217
> mean(y)
[1] 0.2364467
```

n	$\hat{g}_n$	95% conf. interval
100	0.234	(0.222 ; 0.245)
1000	0.236	(0.233 ; 0.240)
10000	0.238	(0.237 ; 0.239)

MC versus MC & Var. red. True value 0.2384217

n	$\hat{g}_n$	95% conf. interval
100	0.206	(0.153 ; 0.259)
1000	0.231	(0.213 ; 0.250)
10000	0.238	(0.232 ; 0.244)

plain MC

n	$\hat{g}_n$	95% conf. interval
100	0.234	(0.222 ; 0.245)
1000	0.236	(0.233 ; 0.240)
10000	0.238	(0.237 ; 0.239)

MC & variance reduction

Control variables

Another approach to variance reduction can be obtained via *control variables*.

Suppose that we can rewrite $\mathbf{E}g(X)$ in the form

$$\mathbf{E}g(X) = \mathbf{E}(g(X) - h(X)) + \mathbf{E}h(X)\,,$$

where $\mathbf{E}h(X)$ can be calculated explicitly and $g(X) - h(X)$ has variance less than $g(X)$.

Then by estimating $\mathbf{E}(g(X) - h(X))$ via the Monte Carlo method, we obtain a reduction in variance.

Let us see another application based on previous example.

Control variables

Consider the price of a *call option*

$$c(X) = \mathbf{E}\left(e^{\beta X} - K\right)_+ .$$

It is easy to show (*put-call parity*) that $c(X) - p(X) = e^{\frac{1}{2}\beta^2} - K$, where p is the price of the put option.

Hence we can write $c(X) = p(X) + e^{\frac{1}{2}\beta^2} - K$.

It is also known that the variance of $p(X)$ is smaller than the variance of $c(X)$. Thus we obtained an estimator of $c(X)$ with reduced bias.

The exact formula for $c(X)$ is also known and reads as

$$\mathbf{E}\left(e^{\beta X} - K\right)_+ = e^{\frac{1}{2}\beta^2}\Phi\left(\beta - \frac{\log(K)}{\beta}\right) - K\Phi\left(-\frac{\log(K)}{\beta}\right) .$$

Control variables: true value 0.887143

```
>  set.seed(123)
>  n <- 1000
>  beta <-1
>  K <- 1
>
>  x <- rnorm(n)
>  y <- sapply(x, function(x) max(0,exp(beta*x)-K))
> mean(y)   # MC value
[1] 0.9030735
>
>  exp(beta^2/2)*pnorm(beta-log(K)/beta)-K*pnorm(-log(K)/beta)
[1] 0.887143  # the true value
>
>  set.seed(123)
>  x <- rexp(n,rate=0.5)
>  h <- function(x) (max(0,1-exp(beta*sqrt(x))) +
+    max(0,1-exp(-beta*sqrt(x))))/sqrt(2*pi*x)
>  w <- sapply(x, h)
>
>  # CALL = PUT + e^{0.5*beta^2} - K
>  z <- w +exp(0.5*beta^2) - K
> mean(z)
[1] 0.885168 # MC & variance reduction
```

MC versus MC & Var. red. True value 0.887143

n	$\hat{g}_n$	95% conf. interval
100	0.858	(0.542 ; 1.174)
1000	0.903	(0.780 ; 1.026)
10000	0.885	(0.844 ; 0.925)

plain MC

n	$\hat{g}_n$	95% conf. interval
100	0.882	(0.871 ; 0.894)
1000	0.885	(0.881 ; 0.889)
10000	0.887	(0.886 ; 0.888)

MC & variance reduction

Antithetic sampling

The idea of antithetic sampling can be applied when it is possible to find transformations of X that leave its distribution unchanged (for example, if X is Gaussian, then $-X$ is Gaussian as well).

Suppose that we want to calculate

$$I = \int_0^1 g(x)\mathrm{d}x = \mathbf{E}g(X)\,,$$

with $X \sim U(0,1)$. The transformation $x \mapsto 1-x$ leaves the distribution unchanged, i.e.

$$1 - X \sim U(0,1)$$

Antithetic sampling

Thus, I can be rewritten as

$$
\begin{aligned}
I &= \frac{1}{2}\int_0^1 (g(x) + g(1-x))\mathrm{d}x \\
&= \frac{1}{2}\mathbf{E}(g(X) + g(1-X)) \\
&= \frac{1}{2}\mathbf{E}(g(X) + g(h(X)))
\end{aligned}
$$

Therefore, we have a variance reduction if

$$
\mathrm{Var}\left(\frac{1}{2}\left(g(X) + g(h(X))\right)\right) < \mathrm{Var}\left(g(X)\right)
$$

Antithetic sampling

So, variance reduction occurs if

$$\mathrm{Var}\left(\frac{1}{2}\left(g(X)+g(h(X))\right)\right) < \mathrm{Var}\left(g(X)\right),$$

which is equivalent to saying that $\mathrm{Cov}(g(X), g(h(X))) < 0$.

Indeed, $\mathrm{Var}(X+Y) = \mathrm{Var}(X) + \mathrm{Var}(Y) + 2\mathrm{Cov}(X, Y)$

Antithetic sampling

So, variance reduction occurs if

$$\mathrm{Var}\left(\frac{1}{2}\left(g(X)+g(h(X))\right)\right)<\mathrm{Var}\left(g(X)\right),$$

which is equivalent to saying that $\mathrm{Cov}(g(X), g(h(X))) < 0$.

Indeed, $\mathrm{Var}(X+Y) = \mathrm{Var}(X) + \mathrm{Var}(Y) + 2\mathrm{Cov}(X, Y)$

But, if $h(x)$ is a monotonic function of x (as in the example above), this is always the case.

This way of proceeding has the effect of reducing the variance but also increasing the accuracy of the calculation of the mean.

Antithetic sampling

So, variance reduction occurs if

$$\mathrm{Var}\left(\frac{1}{2}\left(g(X)+g(h(X))\right)\right) < \mathrm{Var}\left(g(X)\right),$$

which is equivalent to saying that $\mathrm{Cov}(g(X), g(h(X))) < 0$.

Indeed, $\mathrm{Var}(X+Y) = \mathrm{Var}(X) + \mathrm{Var}(Y) + 2\mathrm{Cov}(X,Y)$

But, if $h(x)$ is a monotonic function of x (as in the example above), this is always the case.

This way of proceeding has the effect of reducing the variance but also increasing the accuracy of the calculation of the mean.

Going back to the example of the calculation of the price of a put option, one should calculate it using X and $-X$ and then averaging as follows

Antithetic sampling: true value 0.2384217

```
>  set.seed(123)
>  n <- 1000
>  beta <-1
>  K <- 1
>  x <- rnorm(n)
>  y1 <- sapply(x, function(x) max(0,K-exp(beta*x)))
>  y2 <- sapply(-x, function(x) max(0,K-exp(beta*x)))
>
>  y <- (y1+y2)/2
> mean(y)  # MC with antithetic sampling
[1] 0.2347266
>
> # the true value
>  K*pnorm(log(K)/beta)-exp(beta^2/2)*pnorm(log(K)/beta-beta)
[1] 0.2384217   # the true value
```

MC versus MC & Ant. Samp. True value 0.2384217

n	$\hat{g}_n$	95% conf. interval
100	0.206	(0.153 ; 0.259)
1000	0.231	(0.213 ; 0.250)
10000	0.238	(0.232 ; 0.244)

plain MC

n	$\hat{g}_n$	95% conf. interval
100	0.226	(0.202 ; 0.250)
1000	0.235	(0.226 ; 0.242)
10000	0.238	(0.235 ; 0.240)

MC & antithetic sampling

Relationship with option pricing

In the Black & Scholes theory of option pricing, we model an underlying asset using a stochastic process S_t called geometric Brownian motion which satisfies the stochastic differential equation

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma S_t \mathrm{d}W_t$$

with volatility σ, interest rate r and drift μ.

Relationship with option pricing

In the Black & Scholes theory of option pricing, we model an underlying asset using a stochastic process S_t called geometric Brownian motion which satisfies the stochastic differential equation

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma S_t \mathrm{d}W_t$$

with volatility σ, interest rate r and drift μ.

The price of a derivative is given by the general formula

$$P_t = e^{-r(T-t)} \mathbf{E}\{f(S_T)\}$$

where $f(\cdot)$ is some payoff function, e.g. $f(x) = \max(S_T - K, 0)$ for a european call option with strike price K, maturity T.

Relationship with option pricing

In the Black & Scholes theory of option pricing, we model an underlying asset using a stochastic process S_t called geometric Brownian motion which satisfies the stochastic differential equation

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma S_t \mathrm{d}W_t$$

with volatility σ, interest rate r and drift μ.

The price of a derivative is given by the general formula

$$P_t = e^{-r(T-t)} \mathbf{E}\{f(S_T)\}$$

where $f(\cdot)$ is some payoff function, e.g. $f(x) = \max(S_T - K, 0)$ for a european call option with strike price K, maturity T.

We need to estimate via Monte Carlo the expected value and, to this end, we need to be able to simulate the values of S_T.

R code: `MC.R`