Module 4: Monte Carlo path simulation

Prof. Mike Giles

mike.giles@maths.ox.ac.uk

Oxford University Mathematical Institute

SDE Path Simulation

In Module 2, looked at the case of European options for which the underlying SDE could be integrated exactly.

Now address the more general case in which the solution to the SDE needs to be approximated because

- the option is path-dependent, and/or
- the SDE is not integrable

This lecture will cover:

- Euler-Maruyama discretisation, weak and strong errors
- improved accuracy for path-dependent options

Euler-Maruyama method

The simplest approximation for the scalar SDE

$$\mathrm{d}S = a(S,t)\;\mathrm{d}t + b(S,t)\;\mathrm{d}W$$

is the forward Euler scheme, which is known as the Euler-Maruyama approximation when applied to SDEs:

$$\widehat{S}_{n+1} = \widehat{S}_n + a(\widehat{S}_n, t_n)\, h + b(\widehat{S}_n, t_n)\, \Delta W_n$$

Here h is the timestep, $\widehat{S}_n$ is the approximation to $S(nh)$ and the ΔW_n are i.i.d. $N(0,h)$ Brownian increments.

Euler-Maruyama method

For ODEs, the forward Euler method has $O(h)$ accuracy, and other more accurate methods would usually be preferred.

However, SDEs are very much harder to approximate so the Euler-Maruyama method is used widely in practice.

Numerical analysis is also very difficult and even the definition of “accuracy” is tricky.

Weak convergence

In finance applications, we are mostly concerned with **weak** errors, the error in the expected payoff due to using a finite timestep h.

For a European payoff $f(S(T))$, the weak error is

$$\mathbb{E}[f(S(T))] - \mathbb{E}[f(\widehat{S}_{T/h})]$$

For a path-dependent option, the weak error is

$$\mathbb{E}[f(S)] - \mathbb{E}[\widehat{f}(\widehat{S})]$$

where $f(S)$ is a function of the entire path $S(t)$, and $\widehat{f}(\widehat{S})$ is a corresponding approximation using the whole discrete path.

Weak convergence

Key theoretical result (Bally and Talay, 1995):

If $p(S)$ is the p.d.f. for $S(T)$ and $\widehat{p}(S)$ is the p.d.f. for $\widehat{S}_{T/h}$ computed using the Euler-Maruyama approximation, then under certain conditions on $a(S,t)$ and $b(S,t)$

$$p(S) - \widehat{p}(S) = O(h)$$

and hence

$$\mathbb{E}[f(S(T))] - \mathbb{E}[f(\widehat{S}_{T/h})] = O(h)$$

(This holds even for digital options with discontinuous payoffs $f(S)$. Earlier theory covered only European options such as put and call options with Lipschitz payoffs.)

Weak convergence

Numerical demonstration: Geometric Brownian Motion

$$\mathrm{d}S = r\,S\,\mathrm{d}t + \sigma\,S\,\mathrm{d}W$$

$r = 0.05,\ \sigma = 0.5,\ T = 1$

European call: $S_0 = 100, K = 110$.

Plot shows weak error versus analytic expectation when using 10^8 paths, and also Monte Carlo error (3 standard deviations)

Weak convergence

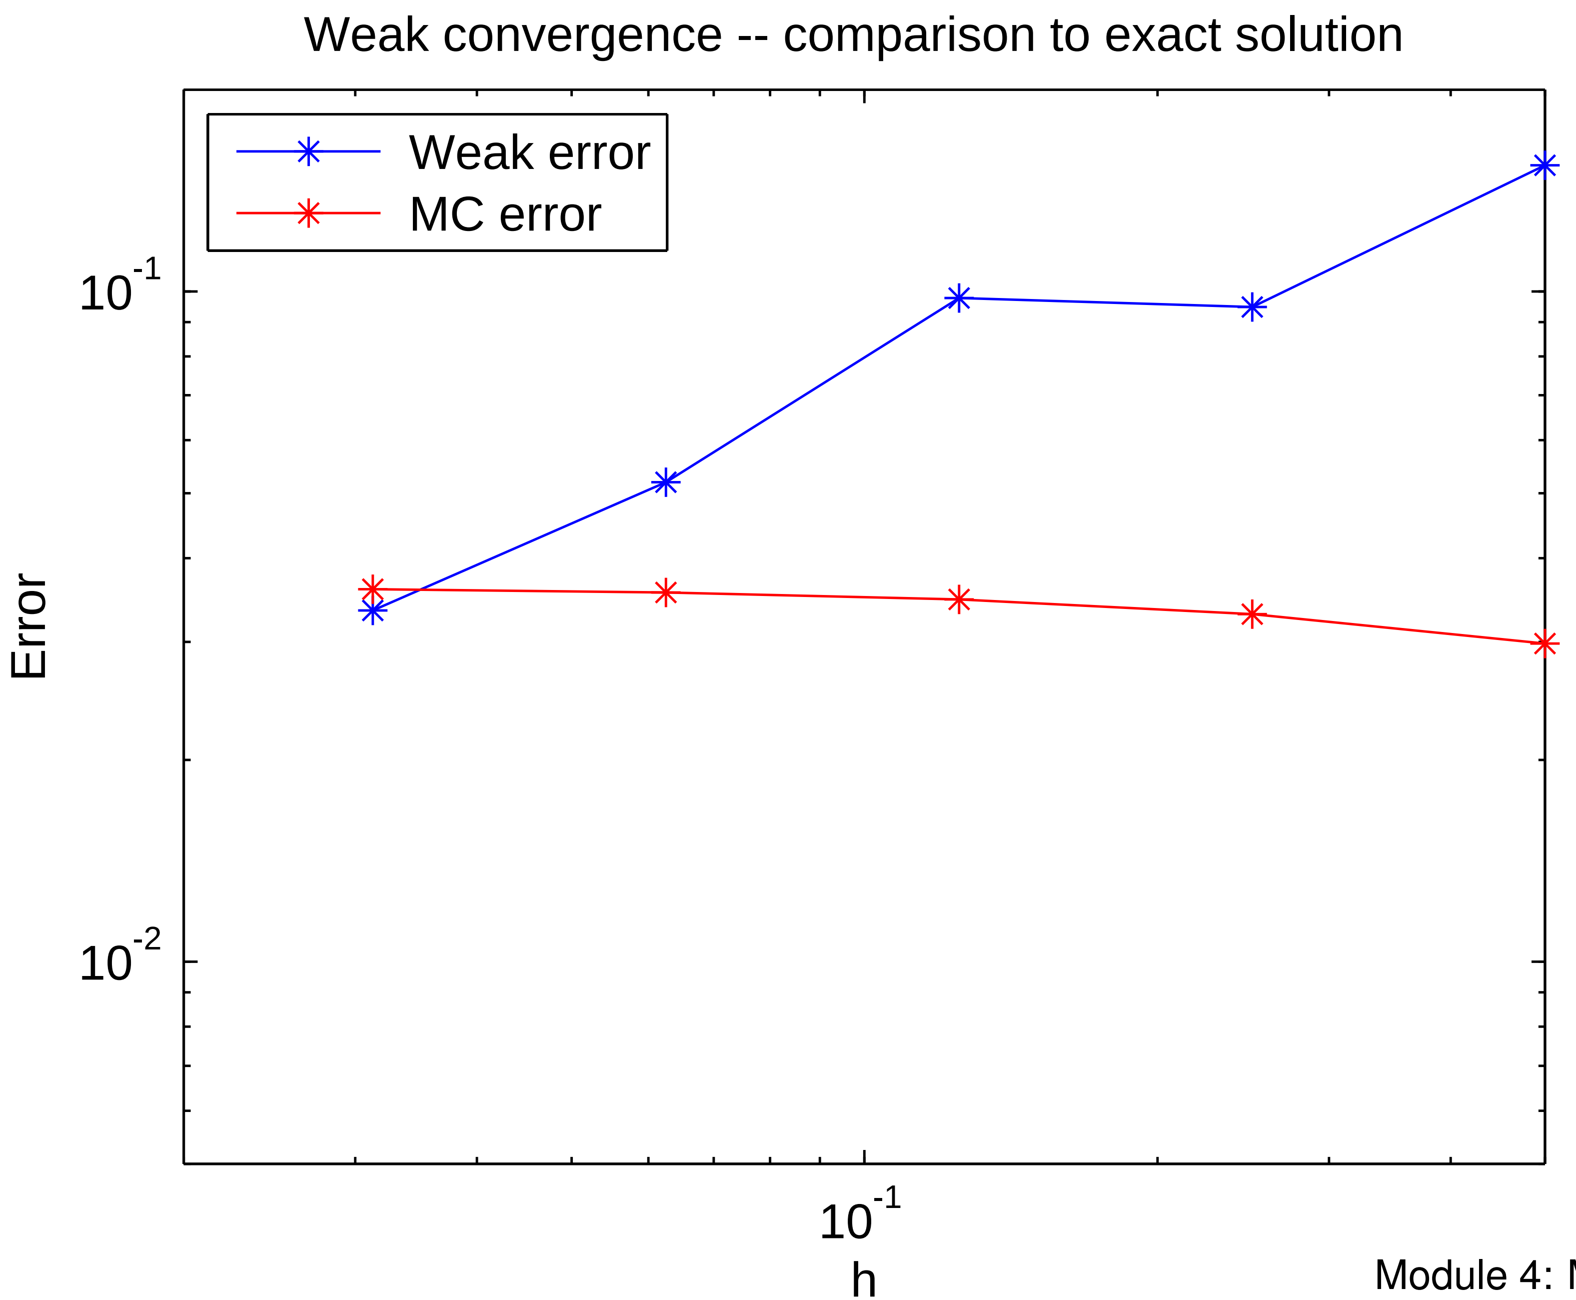

Weak convergence

Previous plot showed difference between exact expectation and numerical approximation.

What if the exact solution is unknown? Compare approximations with timesteps h and $2h$.

If

$$\mathbb{E}[f(S(T))] - \mathbb{E}[f(\widehat{S}^{h}_{T/h})] \approx a\ h$$

then

$$\mathbb{E}[f(S(T))] - \mathbb{E}[f(\widehat{S}^{2h}_{T/2h})] \approx 2\,a\,h$$

and so

$$\mathbb{E}[f(\widehat{S}^{h}_{T/h})] - \mathbb{E}[f(\widehat{S}^{2h}_{T/2h})] \approx a\,h$$

Weak convergence

To minimise the number of paths that need to be simulated, best to use **same** driving Brownian path when doing $2h$ and h approximations – i.e. take Brownian increments for h simulation and sum in pairs to get Brownian increments for $2h$ simulation.

This is like using the same driving Brownian paths for finite difference Greeks. The variance is lower because the h and $2h$ paths are close to each other (**strong** convergence).

(In Module 6, I'll explainhow this forms the basis for the **Multilevel Monte Carlo** method (Giles, 2006))

Weak convergence

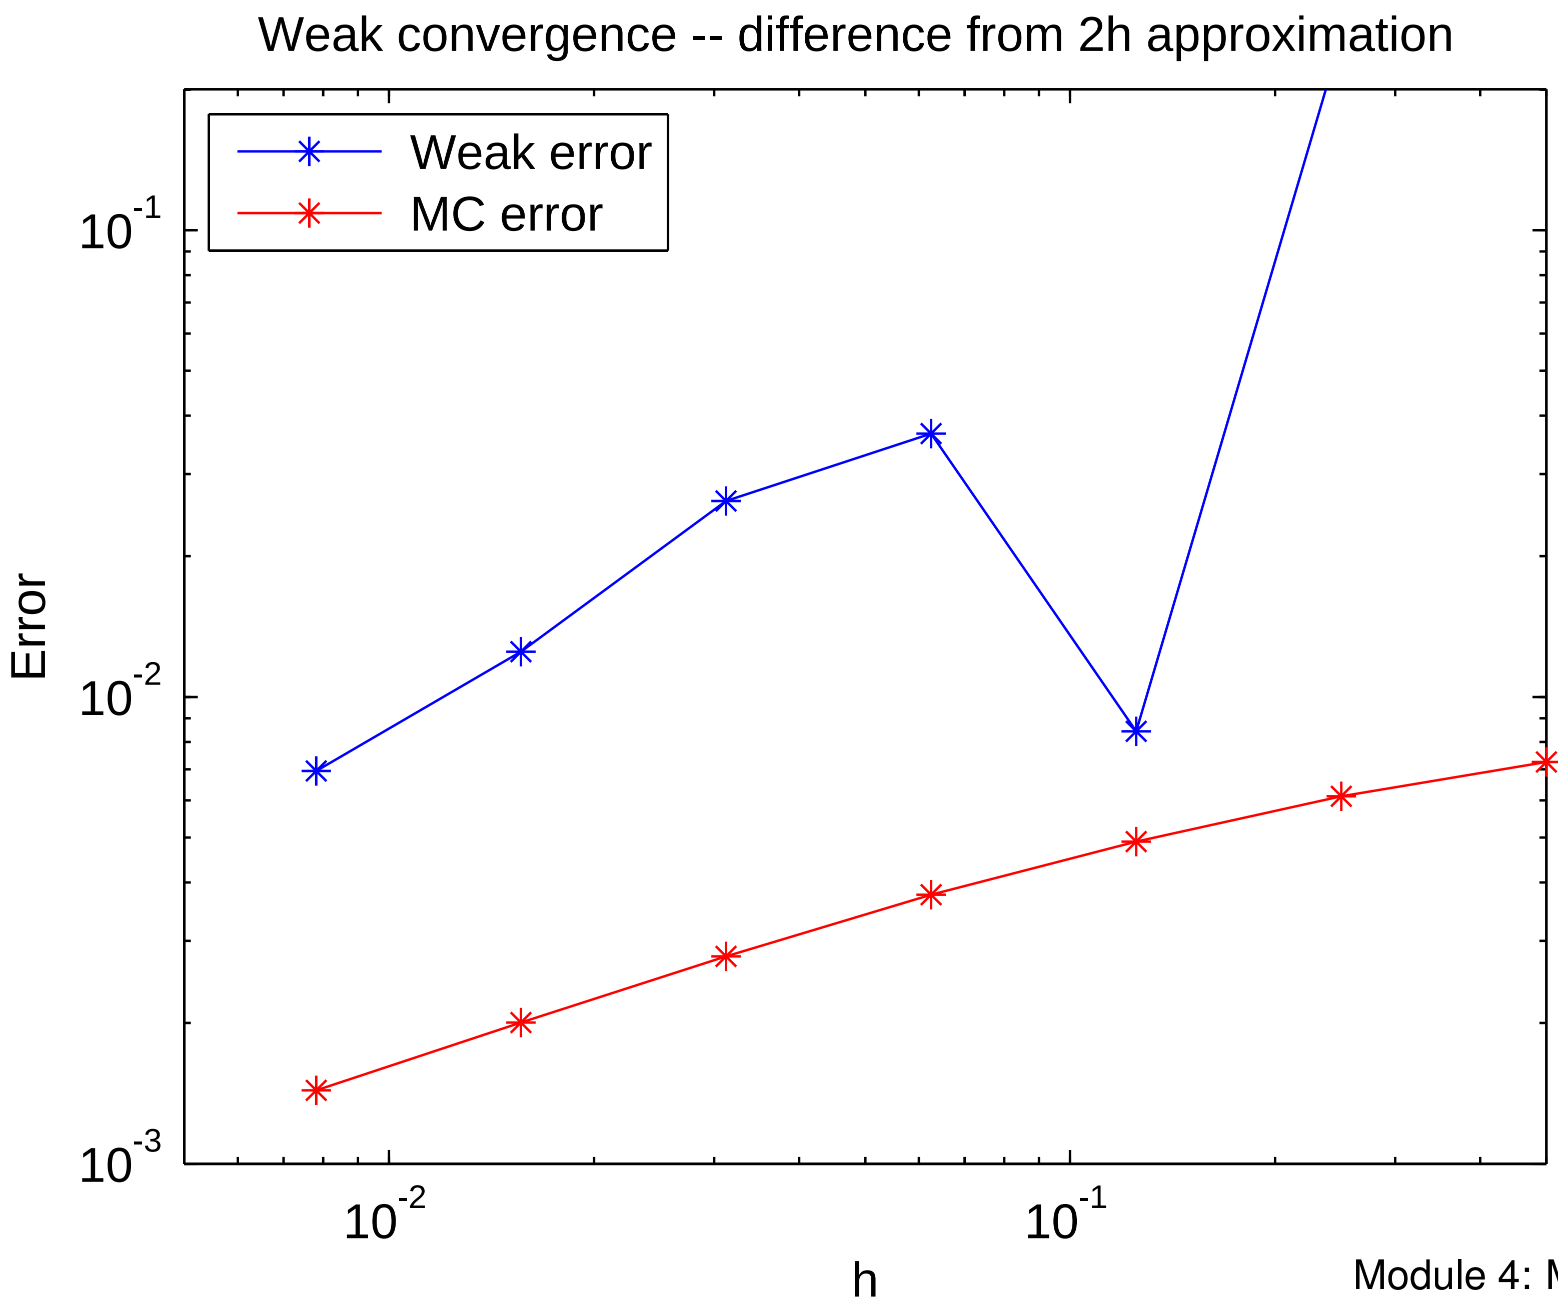

Strong convergence

Strong convergence looks instead at the average error in each individual path:

$$\mathbb{E}\left[\, \left| S(T) - \widehat{S}_{T/h} \right| \,\right] \quad \text{or} \quad \left(\mathbb{E}\left[\left(S(T) - \widehat{S}_{T/h} \right)^2 \right] \right)^{1/2}$$

The main theoretical result (Kloeden & Platen 1992) is that for the Euler-Maruyama method under certain conditions on $a(S,t)$ and $b(S,t)$ these are both $O(\sqrt{h})$.

Strong convergence

Thus, each approximate path deviates by $O(\sqrt{h})$ from its true path.

How can the weak error be $O(h)$? Because the error

$$S(T) - \widehat{S}_{T/h}$$

has mean $O(h)$ even though the r.m.s. is $O(\sqrt{h})$.

(In fact to leading order it is normally distributed with zero mean and variance $O(h)$.)

Strong convergence

Numerical demonstration based on same Geometric Brownian Motion.

Plot shows two curves, one showing the difference from the true solution

$$S(T) = S_0 \, \exp\left((r - \tfrac{1}{2}\sigma^2)T + \sigma\, W(T)\right)$$

and the other showing the difference from the $2h$ approximation

Strong convergence

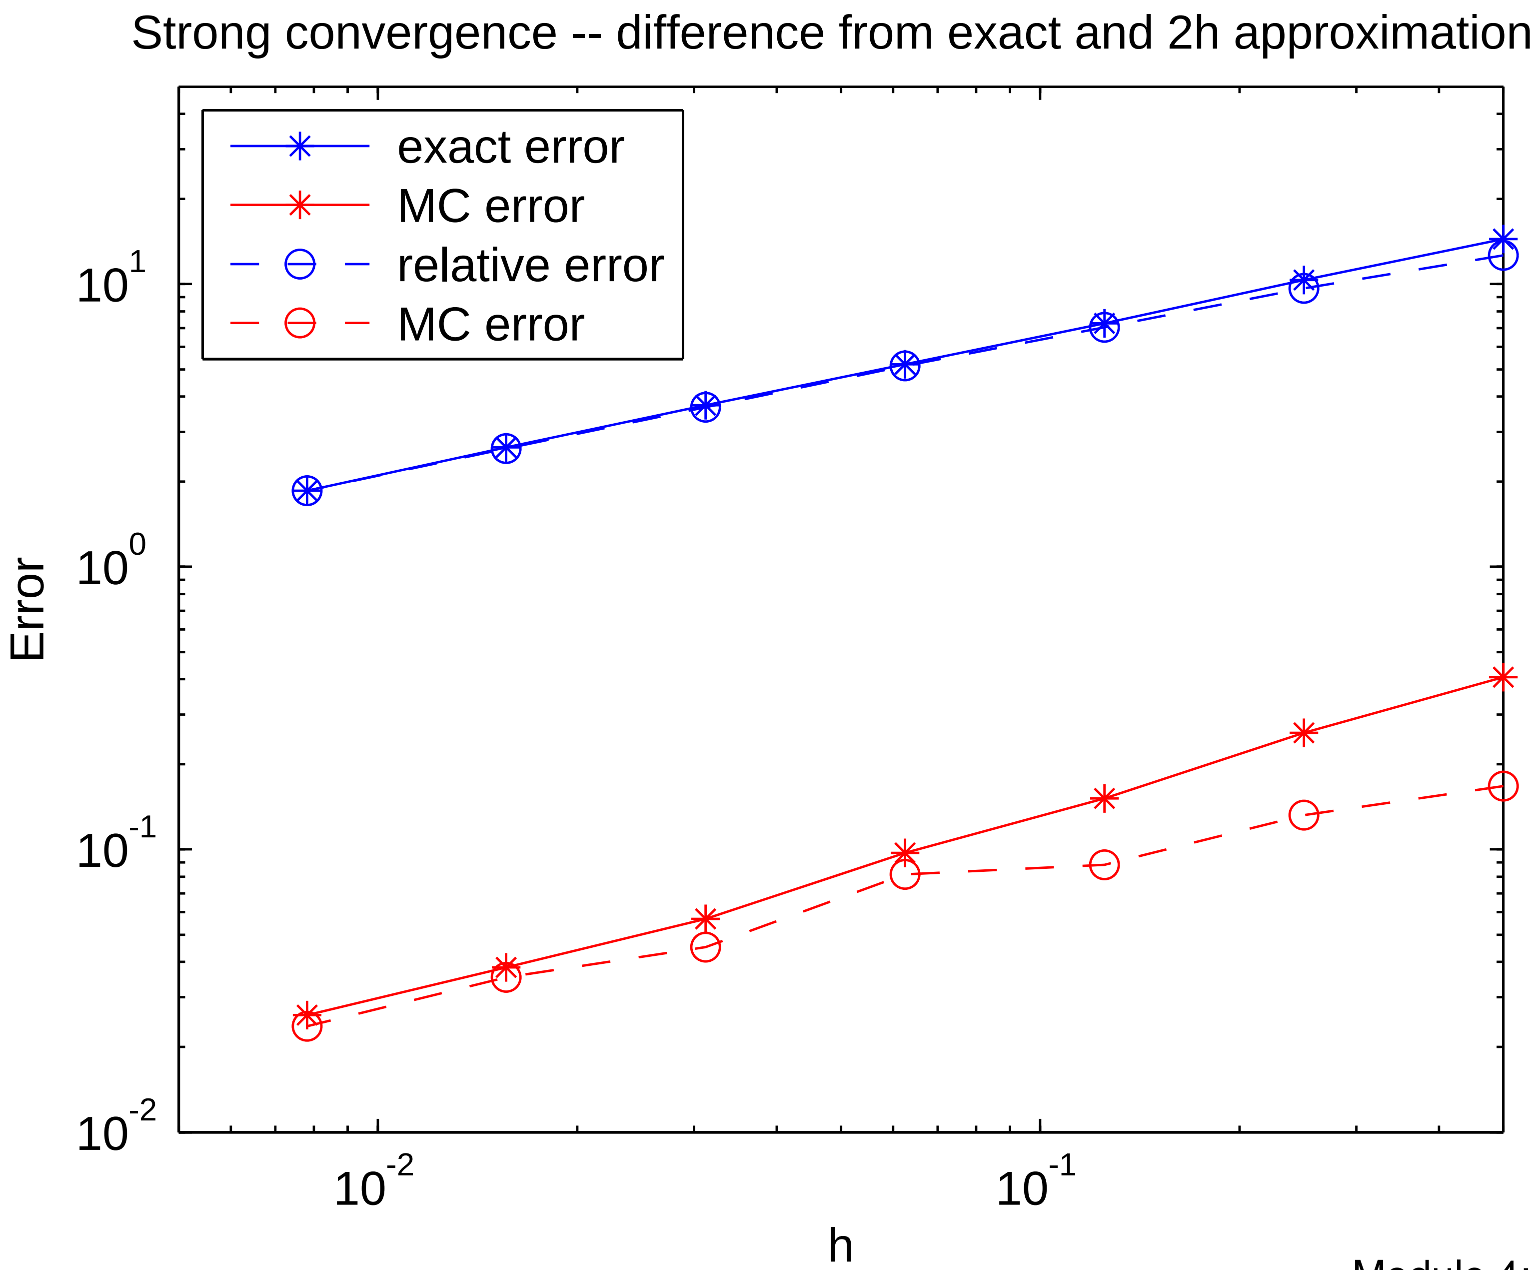

Mean Square Error

Finally, how to decide whether it is better to increase the number of timesteps (reducing the weak error) or the number of paths (reducing the Monte Carlo sampling error)?

If the true option value is $$V = \mathbb{E}[f]$$

and the discrete approximation is $$\widehat{V} = \mathbb{E}[\widehat{f}]$$

and the Monte Carlo estimate is $$\widehat{Y} = \frac{1}{N} \sum_{n=1}^{N} \widehat{f}^{(n)}$$

then ...

Mean Square Error

...the Mean Square Error is

$$\begin{aligned}
\mathbb{E}\left[\left(\widehat{Y}-V\right)^2\right] &= \mathbb{E}\left[\left(\widehat{Y}-\mathbb{E}[\widehat{f}] \;+\; \mathbb{E}[\widehat{f}]-\mathbb{E}[f]\right)^2\right] \\
&= \mathbb{E}\left[(\widehat{Y}-\mathbb{E}[\widehat{f}])^2\right] + (\mathbb{E}[\widehat{f}]-\mathbb{E}[f])^2 \\
&= N^{-1}\mathbb{V}[\widehat{f}] + \left(\mathbb{E}[\widehat{f}]-\mathbb{E}[f]\right)^2
\end{aligned}$$

- first term is due to the variance of estimator
- second term is square of bias due to weak error

Mean Square Error

If there are M timesteps, the computational cost is proportional to $C = NM$ and the MSE is approximately

$$a\,N^{-1} + b\,M^{-2} = a\,N^{-1} + b\,C^{-2}N^2.$$

For a fixed computational cost, this is a minimum when

$$N = \left(\frac{a\,C^2}{2\,b}\right)^{1/3}, \quad M = \left(\frac{2\,b\,C}{a}\right)^{1/3},$$

and hence

$$a\,N^{-1} = \left(\frac{2\,a^2 b}{C^2}\right)^{1/3}, \quad b\,M^{-2} = \left(\frac{a^2 b}{4\,C^2}\right)^{1/3},$$

so the MC term is twice as big as the bias term.

Summary

- simple Euler-Maruyama method is basis for most Monte Carlo simulation in industry – $O(h)$ weak convergence and $O(\sqrt{h})$ strong convergence
- weak convergence is very important when estimating expectations
- strong convergence is usually not important
- Mean-Square-Error is minimised by balancing bias due to weak error and Monte Carlo sampling error

Path-dependent options

For European options, Euler-Maruyama method has $O(h)$ weak convergence.

However, for some path-dependent options it can give only $O(\sqrt{h})$ weak convergence, unless the numerical payoff is constructed carefully.

Barrier option

A down-and-out call option has discounted payoff

$$\exp(-rT)\ (S(T)-K)^{+}\mathbf{1}_{\min_t S(t)>B}$$

i.e. it is like a standard call option except that it pays nothing if the minimum value drops below the barrier B.

The natural numerical discretisation of this is

$$f = \exp(-rT)\ (\widehat{S}_{T/h}-K)^{+}\mathbf{1}_{\min_n \widehat{S}_n>B}$$

Barrier option

Numerical demonstration: Geometric Brownian Motion

$$\mathrm{d}S = r\,S\,\mathrm{d}t + \sigma\,S\,\mathrm{d}W$$

$r = 0.05,\ \sigma = 0.5,\ T = 1$

Down-and-out call: $S_0 = 100, K = 110, B = 90$.

Plots shows weak error versus analytic expectation using 10^6 paths, and difference from $2h$ approximation using 10^5 paths.

(We don’t need as many paths as before because the weak errors are much larger in this case.)

Barrier option

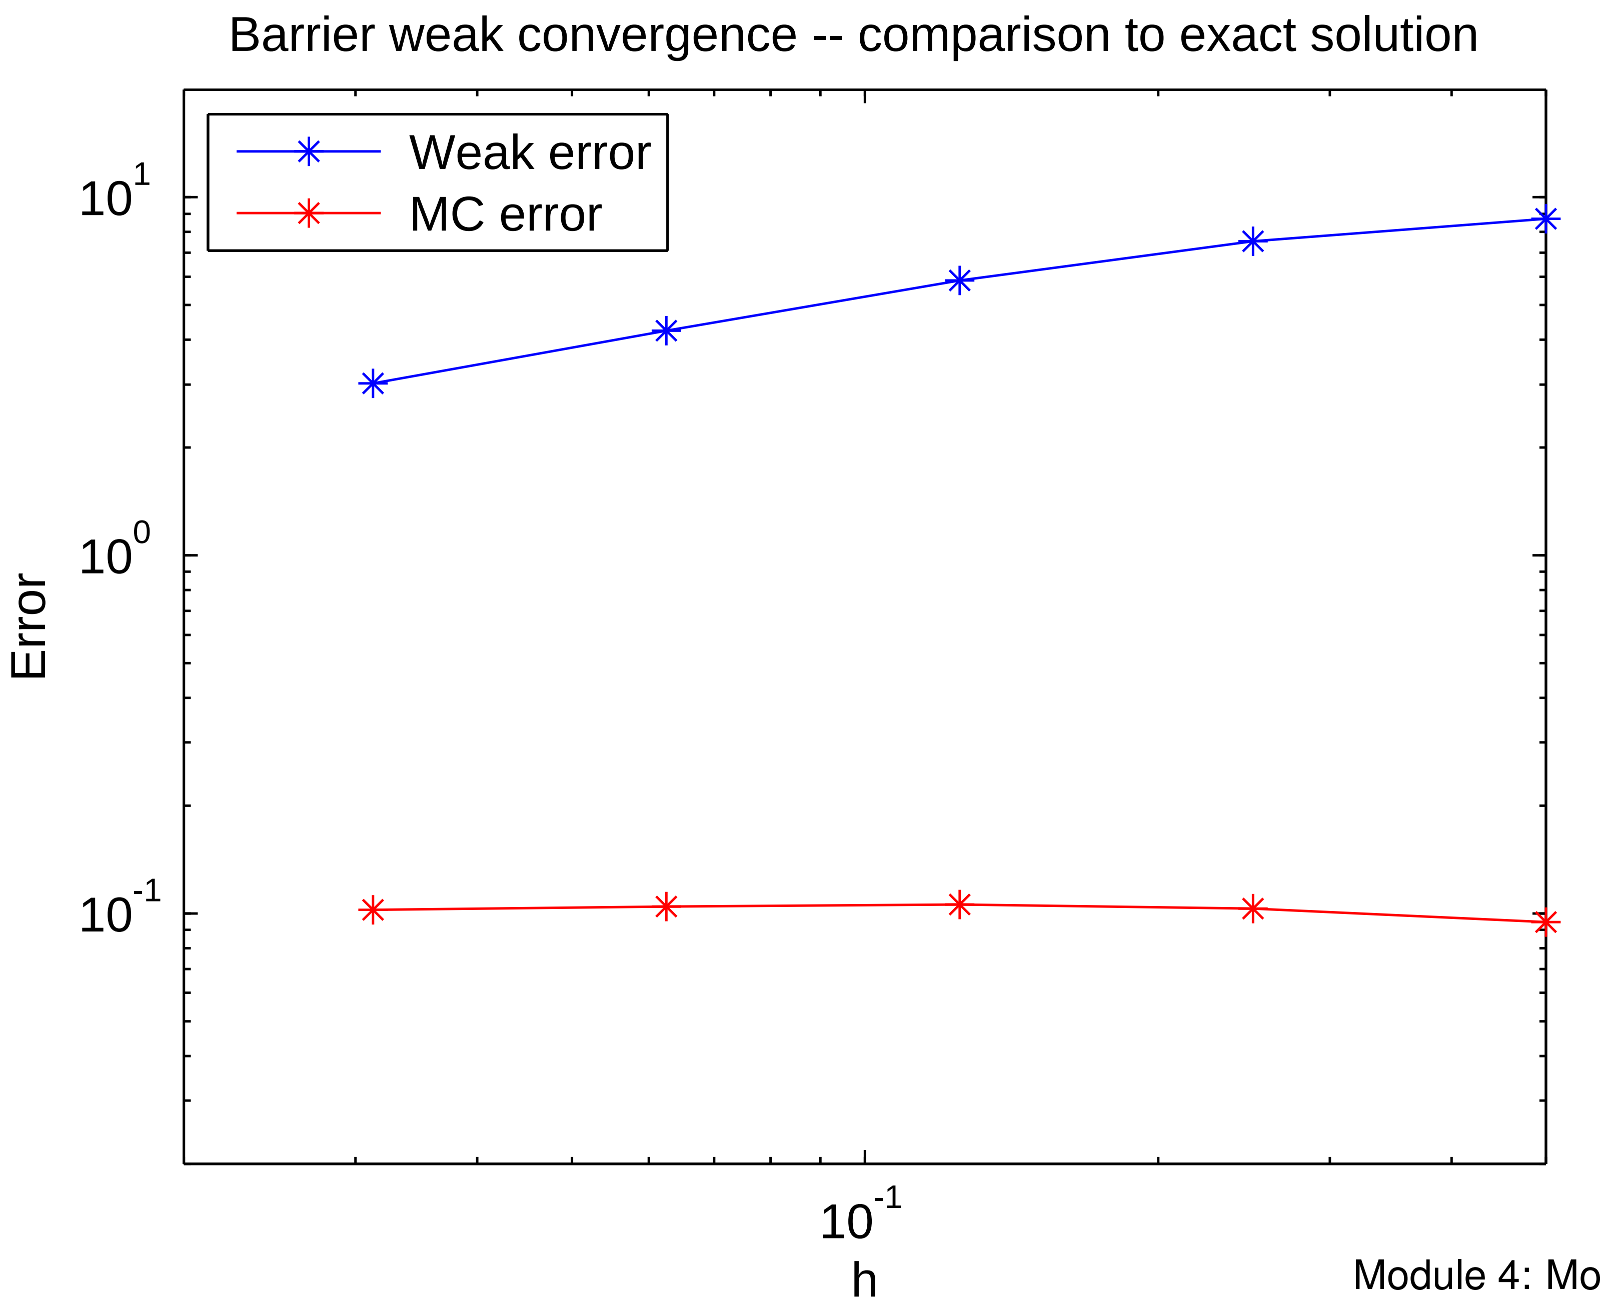

Barrier option

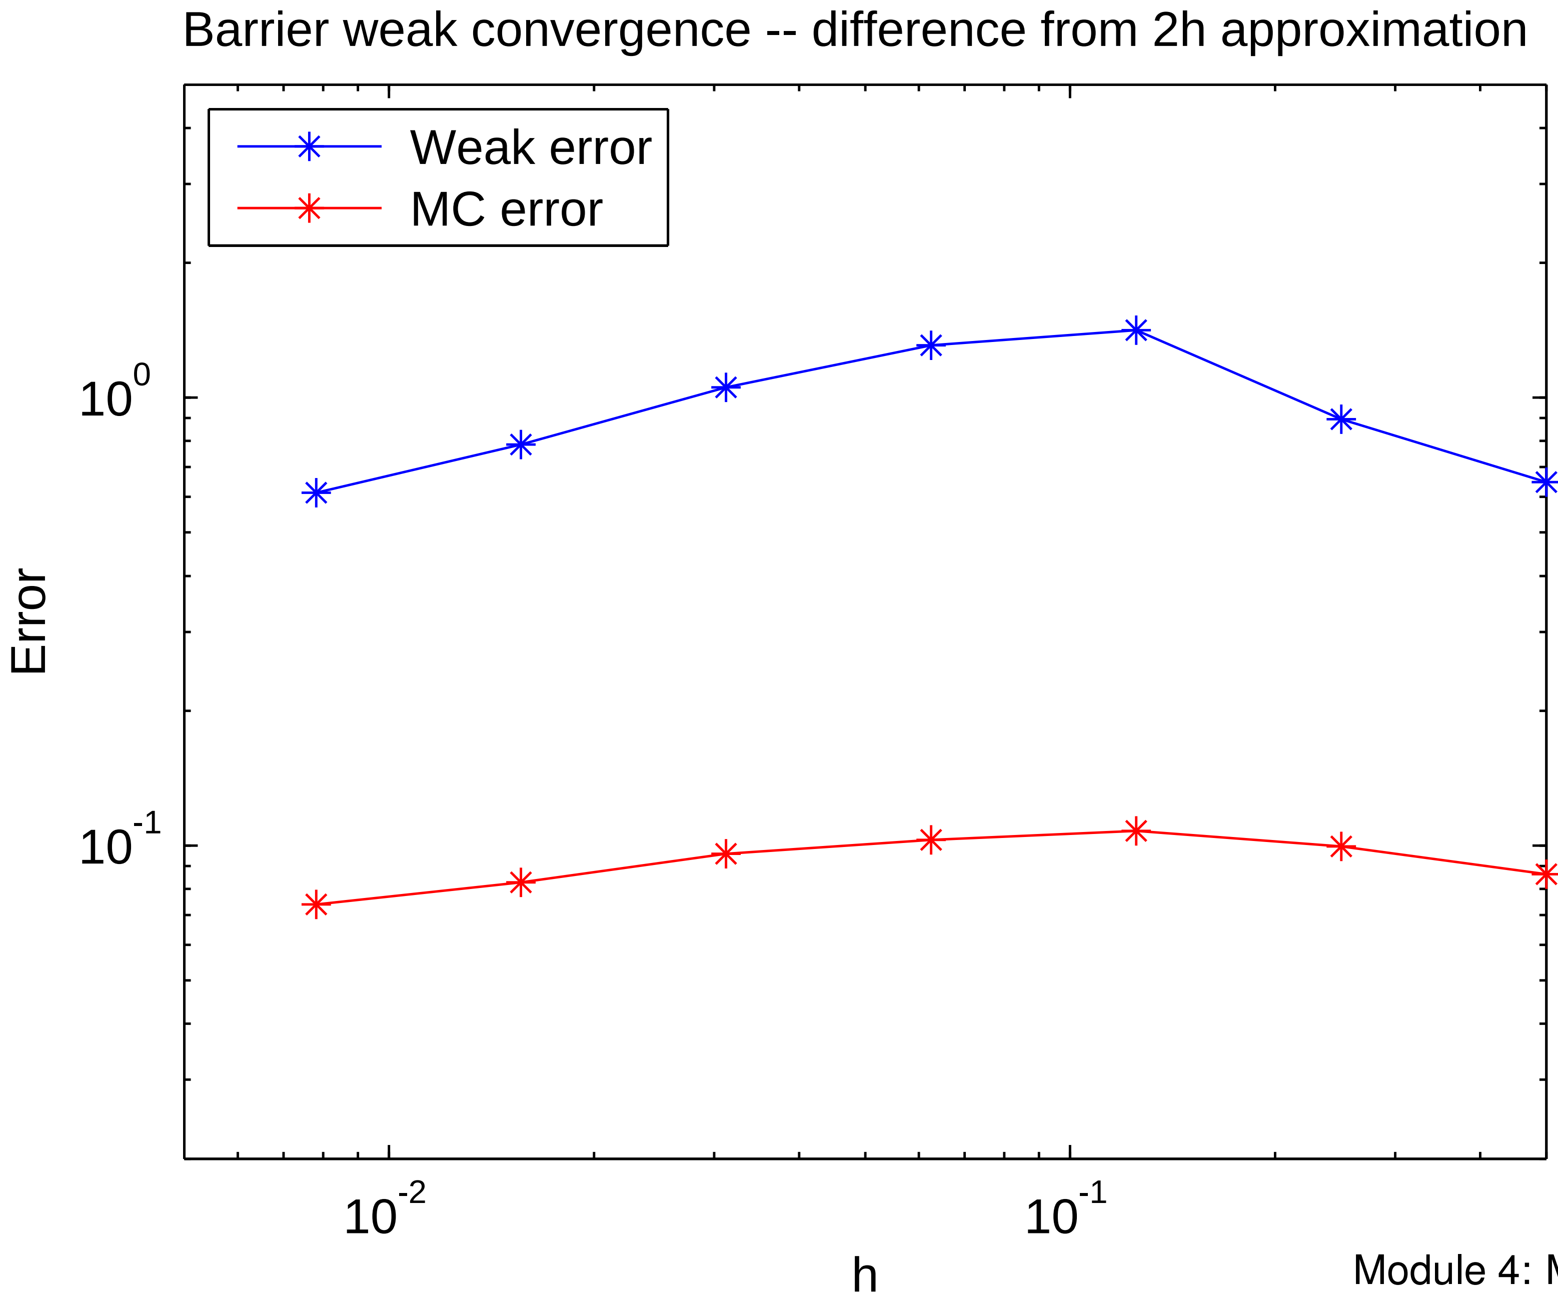

Lookback option

A floating-strike lookback call option has discounted payoff

$$\exp(-rT)\ \left(S(T) - \min_{[0,T]} S(t) \right)$$

The natural numerical discretisation of this is

$$f = \exp(-rT)\ \left(\widehat{S}_{T/h} - \min_n \widehat{S}_n \right)$$

Lookback option

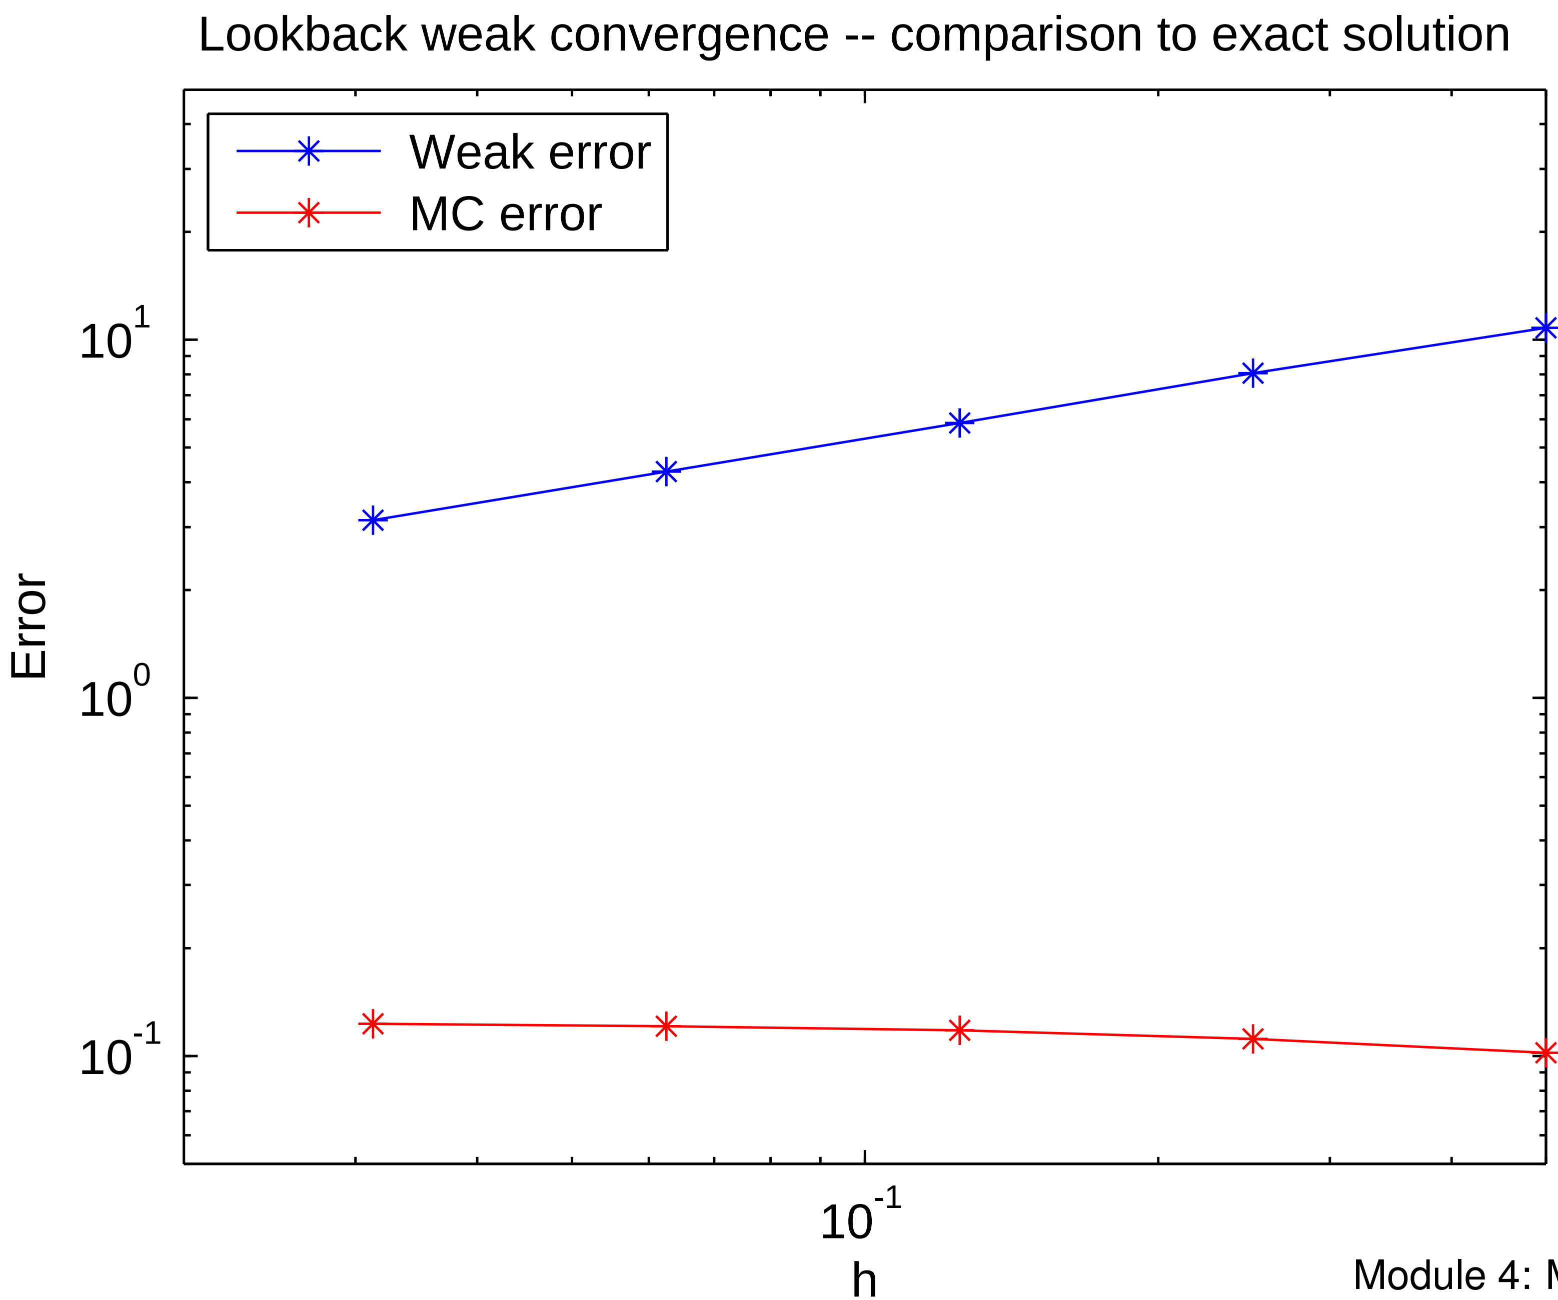

Lookback option

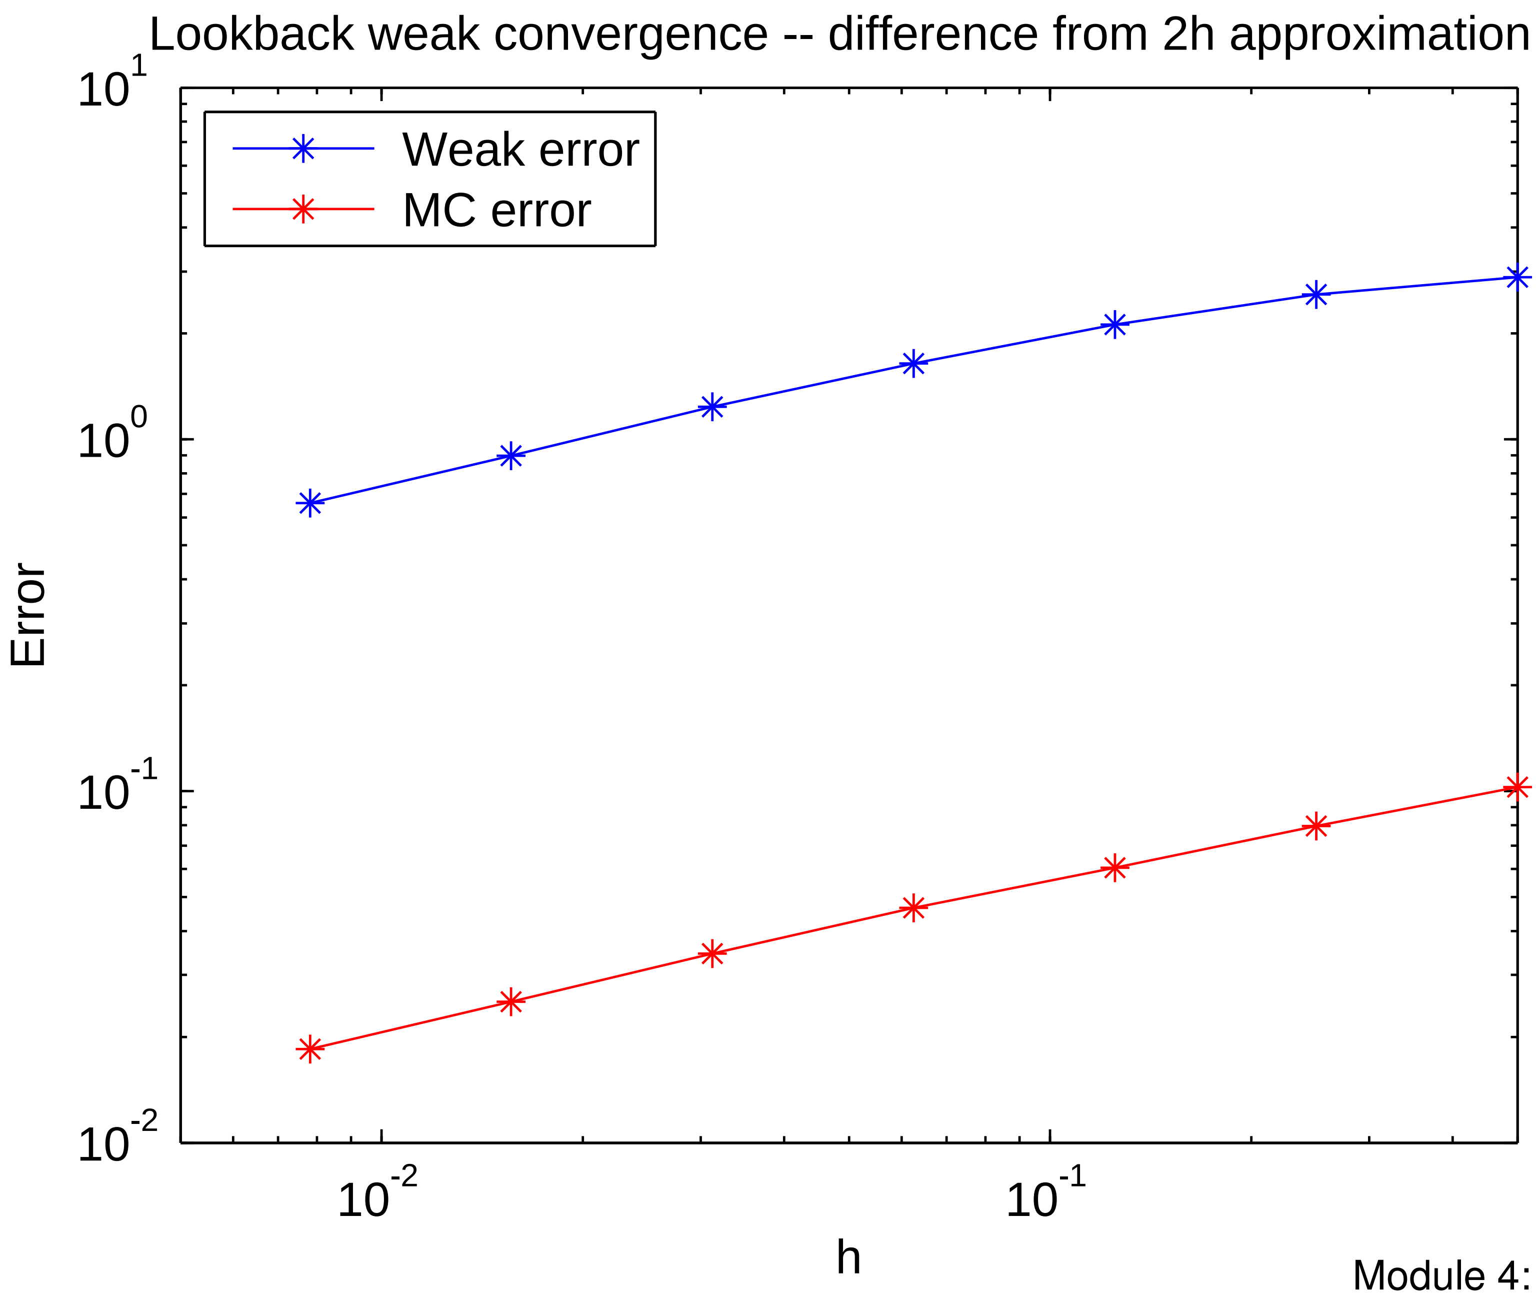

Brownian bridge

To recover $O(h)$ weak convergence we first need some theory.

Consider simple Brownian motion

$$\mathrm{d}S = a\ \mathrm{d}t + b\ \mathrm{d}W$$

with constant $a,\ b$ and initial data $S(0)\!=\!0$.

Question: given $S(T)$, what is conditional probability density for $S(T/2)$?

Conditional probability

With discrete probabilities,

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

Similarly, with probability density functions

$$p_1(x|y) = \frac{p_2(x,y)}{p_3(y)}$$

where

- $p_1(x|y)$ is the conditional p.d.f. for x, given y
- $p_2(x,y)$ is the joint probability density function for x, y
- $p_3(y)$ is the probability density function for y

Brownian bridge

In our case,

$$y \equiv S(T), \quad x \equiv S(T/2)$$

$$\begin{aligned}
p_2(x,y) &= \frac{1}{\sqrt{\pi\, T}\; b}\; \exp\left(-\frac{(x-aT/2)^2}{b^2\, T}\right) \\
&\times \frac{1}{\sqrt{\pi\, T}\; b}\; \exp\left(-\frac{(y-x-aT/2)^2}{b^2\, T}\right) \\
p_3(y) &= \frac{1}{\sqrt{2\pi\, T}\; b}\; \exp\left(-\frac{(y-aT)^2}{2\, b^2\, T}\right) \\
\Longrightarrow p_1(x|y) &= \frac{1}{\sqrt{\pi\, T/2}\; b}\; \exp\left(-\frac{(x-y/2)^2}{b^2\, T/2}\right)
\end{aligned}$$

Hence, x is Normally distributed with mean $y/2$ and variance $b^2 T/4$.

Brownian bridge

Extending this to a particular timestep with endpoints $S(t_n)$ and $S(t_{n+1})$, conditional on these the mid-point is Normally distributed with mean

$$\frac{1}{2}\left(S(t_n) + S(t_{n+1})\right)$$

and variance $b^2 h/4$.

We can take a sample from this conditional p.d.f. and then repeat the process, recursively bisecting each interval to fill in more and more detail.

Note: the drift a is irrelevant, given the two endpoints. Because of this, we will take $a=0$ in the next bit of theory.

Barrier crossing

Consider zero drift Brownian motion with $S(0)\!>\!0$.

If the path $S(t)$ hits a barrier at 0, it is equally likely thereafter to go up or down. Hence, by symmetry, for $s>0$, the p.d.f. for paths with $S(T)=s$ after hitting the barrier is equal to the p.d.f. for paths with $S(T)=-s$.

Thus, for $S(T)>0$,

$$\begin{aligned} P(\text{hit barrier}|S(T)) &= \frac{\exp\left(-\frac{(-S(T)-S(0))^2}{2b^2T}\right)}{\exp\left(-\frac{(S(T)-S(0))^2}{2b^2T}\right)} \\ &= \exp\left(-\frac{2\,S(T)\,S(0)}{b^2T}\right) \end{aligned}$$

Barrier crossing

For a timestep $[t_n, t_{n+1}]$ and non-zero barrier B this generalises to

$$P(\text{hit barrier}|S_n, S_{n+1} > B) = \exp\left(- \frac{2\,(S_{n+1}-B)\,(S_n-B)}{b^2 h} \right)$$

This can also be viewed as the cumulative probability $P(S_{min} < B)$ where $S_{min} = \min\limits_{[t_n, t_{n+1}]} S(t)$.

Since this is uniformly distributed on $[0,1]$ we can equate this to a uniform $[0,1]$ random variable U_n and solve to get

$$S_{min} = \tfrac{1}{2}\left(S_{n+1} + S_n - \sqrt{(S_{n+1}-S_n)^2 - 2\,b^2 h\, \log U_n} \right)$$

Barrier crossing

For a barrier above, we have

$$P(\text{hit barrier}|S_n, S_{n+1} < B) = \exp\left(-\,\frac{2\,(B-S_{n+1})\,(B-S_n)}{b^2 h} \right)$$

and hence

$$S_{max} = \tfrac{1}{2}\left(S_{n+1} + S_n + \sqrt{(S_{n+1}-S_n)^2 - 2\,b^2 h\ \log U_n} \right)$$

where U_n is again a uniform $[0,1]$ random variable.

Barrier option

Returning now to the barrier option, how do we define the numerical payoff $\widehat{f}(\widehat{S})$?

First, calculate $\widehat{S}_n$ as usual using Euler-Maruyama method.

Second, two alternatives:

- use (approximate) probability of crossing the barrier
- directly sample (approximately) the minimum in each timestep

Barrier option

Alternative 1: treating the drift and volatility as being approximately constant within each timestep, the probability of having crossed the barrier within timestep n is

$$P_n = \exp\left(- \frac{2\,(\widehat{S}_{n+1} - B)^+\,(\widehat{S}_n - B)^+}{b^2(\widehat{S}_n, t_n)\; h} \right)$$

Probability at end of not having crossed barrier is $\prod_n (1 - P_n)$ and so the payoff is

$$\widehat{f}(\widehat{S}) = \exp(-rT)\;(\widehat{S}_{T/h} - K)^+\;\prod_n (1 - P_n).$$

I prefer this approach because it is differentiable – good for Greeks

Barrier option

Alternative 2: again treating the drift and volatility as being approximately constant within each timestep, define the minimum within timestep n as

$$\widehat{M}_n = \tfrac{1}{2}\left(\widehat{S}_{n+1} + \widehat{S}_n - \sqrt{(\widehat{S}_{n+1} - \widehat{S}_n)^2 - 2\, b^2(\widehat{S}_n, t_n)\, h\, \log U_n} \right)$$

where the U_n are i.i.d. uniform $[0,1]$ random variables.

The payoff is then

$$\widehat{f}(\widehat{S}) = \exp(-rT)\ (\widehat{S}_{T/h} - K)^+\ \mathbf{1}_{\min_n \widehat{M}_n > B}$$

With this approach one can stop the path calculation as soon as one $\widehat{M}_n$ drops below B.

Lookback option

This is treated in a similar way to Alternative 2 for the barrier option.

We construct a minimum $\widehat{M}_n$ within each timestep and then the payoff is

$$\widehat{f}(\widehat{S}) = \exp(-rT)\ \left(\widehat{S}_{T/h} - \min_n \widehat{M}_n\right)$$

This is differentiable, so good for Greeks – unlike Alternative 2 for the barrier option.

Weak convergence

With these modification to the numerical payoff approximation, the weak convergence for both barrier and lookback options is improved from $O(\sqrt{h})$ to $O(h)$.

See practical for numerical demonstration!

Final Words

- "natural" approximation of barrier and lookback options leads to poor $O(\sqrt{h})$ weak convergence
- this is an inevitable consequence of dependence on minimum/maximum and $O(\sqrt{h})$ path variation within each timestep
- improved treatment based on Brownian bridge theory approximates behaviour within timestep as simple Brownian motion with constant drift and volatility
  – gives $O(h)$ weak convergence