

Monte Carlo Methods

Prof. Mike Giles

`mike.giles@maths.ox.ac.uk`

Oxford University Mathematical Institute

Greeks

In Monte Carlo applications we don't just want to know the expected discounted value of some payoff

$$V = \mathbb{E}[f(S(T))]$$

For hedging and risk analysis, we also want to know a whole range of “Greeks” corresponding to first and second derivatives of V with respect to various parameters:

$$\Delta = \frac{\partial V}{\partial S_0}, \quad \Gamma = \frac{\partial^2 V}{\partial S_0^2},$$
$$\rho = \frac{\partial V}{\partial r}, \quad \text{Vega} = \frac{\partial V}{\partial \sigma}.$$

Finite difference sensitivities

If $V(\theta) = \mathbb{E}[f(S(T))]$ for an input parameter θ is sufficiently differentiable, then the sensitivity $\frac{\partial V}{\partial \theta}$ can be approximated by one-sided finite difference

$$\frac{\partial V}{\partial \theta} = \frac{V(\theta + \Delta\theta) - V(\theta)}{\Delta\theta} + O(\Delta\theta)$$

or by central finite difference

$$\frac{\partial V}{\partial \theta} = \frac{V(\theta + \Delta\theta) - V(\theta - \Delta\theta)}{2\Delta\theta} + O((\Delta\theta)^2)$$

(This approach is referred to as getting Greeks by “bumping” the input parameters.)

Finite difference sensitivities

The clear advantage of this approach is that it is very simple to implement (hence the most popular in practice?)

However, the disadvantages are:

- expensive (2 extra sets of calculations for central differences)
- significant bias error if $\Delta\theta$ too large
- machine roundoff errors if $\Delta\theta$ too small
- large variance if $f(S(T))$ discontinuous and $\Delta\theta$ small

Finite difference sensitivities

Let $X^{(i)}(\theta + \Delta\theta)$ and $X^{(i)}(\theta - \Delta\theta)$ be the values of $f(S(T))$ obtained for different MC samples, so the central difference estimate for $\frac{\partial V}{\partial \theta}$ is given by

$$\begin{aligned}\hat{Y} &= \frac{1}{2\Delta\theta} \left(N^{-1} \sum_{i=1}^N X^{(i)}(\theta + \Delta\theta) - N^{-1} \sum_{i=1}^N X^{(i)}(\theta - \Delta\theta) \right) \\ &= \frac{1}{2N\Delta\theta} \sum_{i=1}^N \left(X^{(i)}(\theta + \Delta\theta) - X^{(i)}(\theta - \Delta\theta) \right)\end{aligned}$$

Finite difference sensitivities

If independent samples are taken for both $X^{(i)}(\theta + \Delta\theta)$ and $X^{(i)}(\theta - \Delta\theta)$ then

$$\begin{aligned}\mathbb{V}[\widehat{Y}] &\approx \left(\frac{1}{2N\Delta\theta}\right)^2 \sum_j \left(\mathbb{V}[X(\theta + \Delta\theta)] + \mathbb{V}[X(\theta - \Delta\theta)]\right) \\ &\approx \left(\frac{1}{2N\Delta\theta}\right)^2 2N \mathbb{V}[f] \\ &= \frac{\mathbb{V}[f]}{2N(\Delta\theta)^2}\end{aligned}$$

which is very large for $\Delta\theta \ll 1$.

Finite difference sensitivities

It is much better for $X^{(i)}(\theta + \Delta\theta)$ and $X^{(i)}(\theta - \Delta\theta)$ to use the same set of random inputs.

If $X^{(i)}(\theta)$ is differentiable with respect to θ , then

$$X^{(i)}(\theta + \Delta\theta) - X^{(i)}(\theta - \Delta\theta) \approx 2 \Delta\theta \frac{\partial X^{(i)}}{\partial \theta}$$

and hence

$$\mathbb{V}[\hat{Y}] \approx N^{-1} \mathbb{V} \left[\frac{\partial X}{\partial \theta} \right],$$

which behaves well for $\Delta\theta \ll 1$, so one should choose a small value for $\Delta\theta$ to minimise the bias due to the finite differencing.

Finite difference sensitivities

However, there are problems if $\Delta\theta$ is chosen to be extremely small.

In finite precision arithmetic, computing $X(\theta)$ has an error which is approximately random with r.m.s. magnitude δ

- single precision $\delta \approx 10^{-6}|X|$
- double precision $\delta \approx 10^{-14}|X|$

Hence, should choose bump $\Delta\theta$ so that

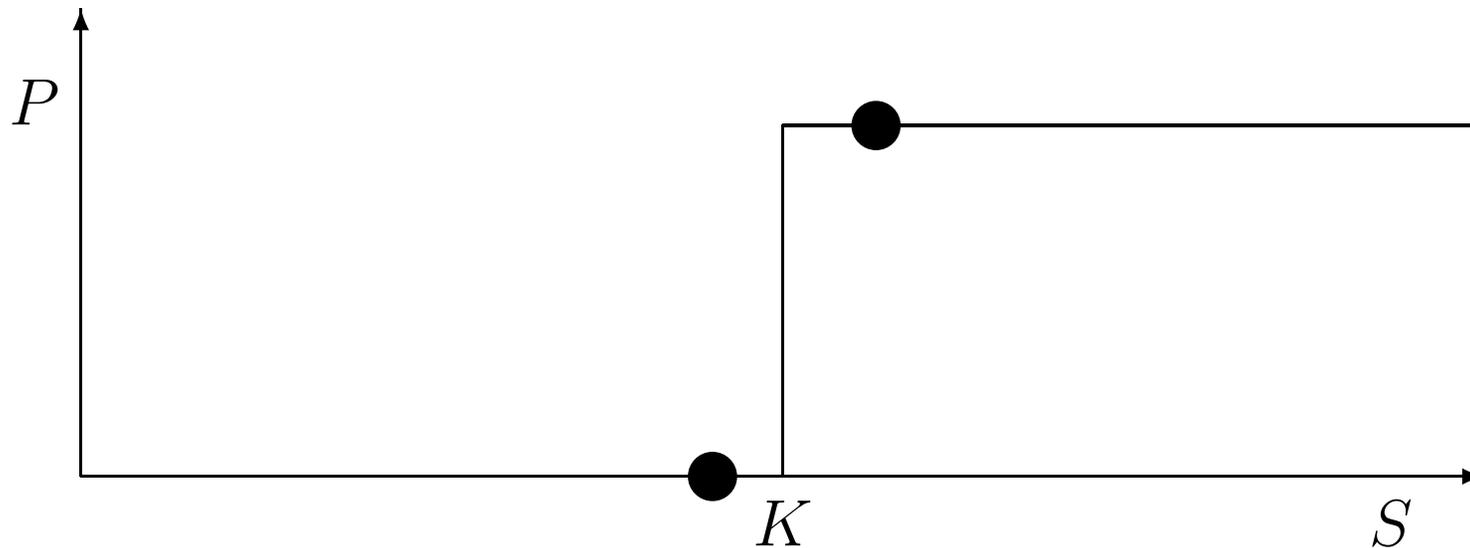
$$|X(\theta + \Delta\theta) - X(\theta - \Delta\theta)| \gg \delta$$

Most banks probably use double precision to be safe.

Finite difference sensitivities

Next, we analyse the variance of the finite difference estimator when the payoff is discontinuous.

The problem is that a small bump in the asset S can produce a big bump in the payoff – not differentiable



Finite difference sensitivities

What is the probability that $S(\theta \pm \Delta\theta)$ will be on different sides of the discontinuity?

Separation of $S(\theta \pm \Delta\theta)$ is $O(\Delta\theta)$

$$\mathbb{P}(|S(\theta) - K| < c \Delta\theta) = O(\Delta\theta)$$

Hence, $O(\Delta\theta)$ probability of straddling the strike.

This leads to a variance which is $O(N^{-1} \Delta\theta^{-1})$.

Finite difference sensitivities

So, small $\Delta\theta$ gives a large variance, while a large $\Delta\theta$ gives a large finite difference discretisation error.

To determine the optimum choice we use this result:

$$\mathbb{E} \left[\left(\hat{Y} - \mathbb{E}[Y] \right)^2 \right] = \mathbb{V}[\hat{Y}] + \left(\mathbb{E}[\hat{Y}] - \mathbb{E}[Y] \right)^2$$

$$\text{Mean Square Error} = \text{variance} + (\text{bias})^2$$

In our case, we have

$$\mathbb{V}[\hat{Y}] + (\text{bias})^2 \sim \frac{a}{N \Delta\theta} + b \Delta\theta^4.$$

which can be minimised by choosing $\Delta\theta$ appropriately.

Pathwise sensitivities

We start with

$$V \equiv \mathbb{E} [f(S(T))] = \int f(S(T)) p_W(W) dW$$

where $p_W(W)$ is the probability density function for $W(T)$, and differentiate this to get

$$\frac{\partial V}{\partial \theta} = \int \frac{\partial f}{\partial S} \frac{\partial S(T)}{\partial \theta} p_W dW = \mathbb{E} \left[\frac{\partial f}{\partial S} \frac{\partial S(T)}{\partial \theta} \right]$$

with $\partial S(T)/\partial \theta$ being evaluated at fixed W .

Pathwise sensitivities

This leads to the estimator

$$\frac{1}{N} \sum_{i=1}^N \frac{\partial f}{\partial S}(S^{(i)}) \frac{\partial S^{(i)}}{\partial \theta}$$

which is the derivative of the usual price estimator

$$\frac{1}{N} \sum_{i=1}^N f(S^{(i)})$$

Gives incorrect estimates when $f(S)$ is discontinuous.

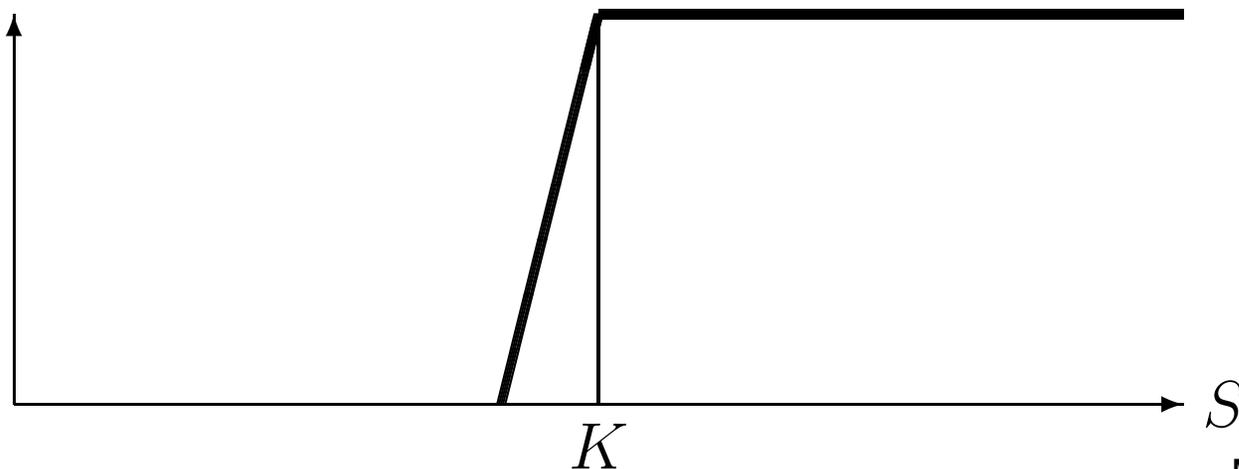
e.g. for digital put $\frac{\partial f}{\partial S} = 0$ so estimated value of Greek is zero – clearly wrong.

Pathwise sensitivities

To handle payoffs which do not have the necessary continuity/smoothness one can modify the payoff

For digital options it is common to use a piecewise linear approximation to limit the magnitude of Δ near maturity
– avoids large transaction costs

Bank selling the option will price it conservatively
(i.e. over-estimate the price)



SDE Path Simulation

So far, we have considered European options with assets satisfying Geometric Brownian Motion SDEs

Now we consider 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

Euler-Maruyama method

The simplest approximation for the scalar SDE

$$dS = a(S, t) dt + b(S, t) dW$$

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

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

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

For ODEs, the $O(h)$ accuracy of forward Euler method is considered poor, but for SDEs it is very hard to do better.

Weak convergence

In finance applications, mostly concerned with **weak** errors, the error in the expected payoff. For a European payoff $f(S(T))$ this is

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

and it is of order α if

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

For a path-dependent option, the weak error is

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

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

Weak convergence

Key theoretical result (Bally and Talay, 1995):

If $p(S)$ is the p.d.f. for $S(T)$ and $\hat{p}(S)$ is the p.d.f. for $\hat{S}_{T/h}$ computed using the Euler-Maruyama approximation, then if $a(S, t)$ and $b(S, t)$ are Lipschitz w.r.t. S, t

$$\|p(S) - \hat{p}(S)\|_1 = O(h)$$

and hence for bounded payoffs

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

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

$$dS = r S dt + \sigma S dW$$

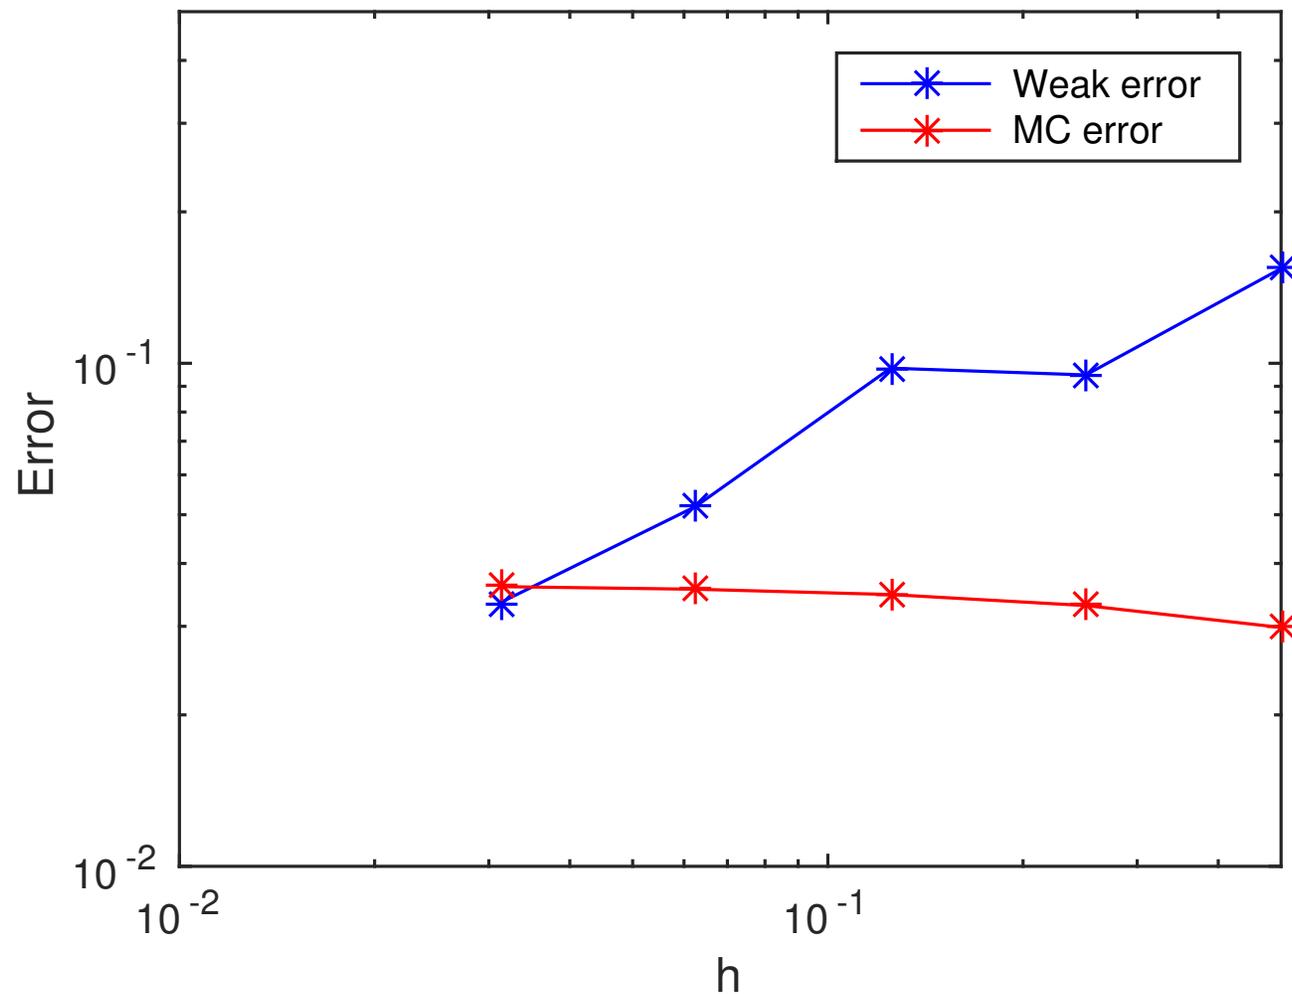
$$r = 0.05, \sigma = 0.5, T = 1$$

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

Comparison to exact solution:



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

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

then

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

and so

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

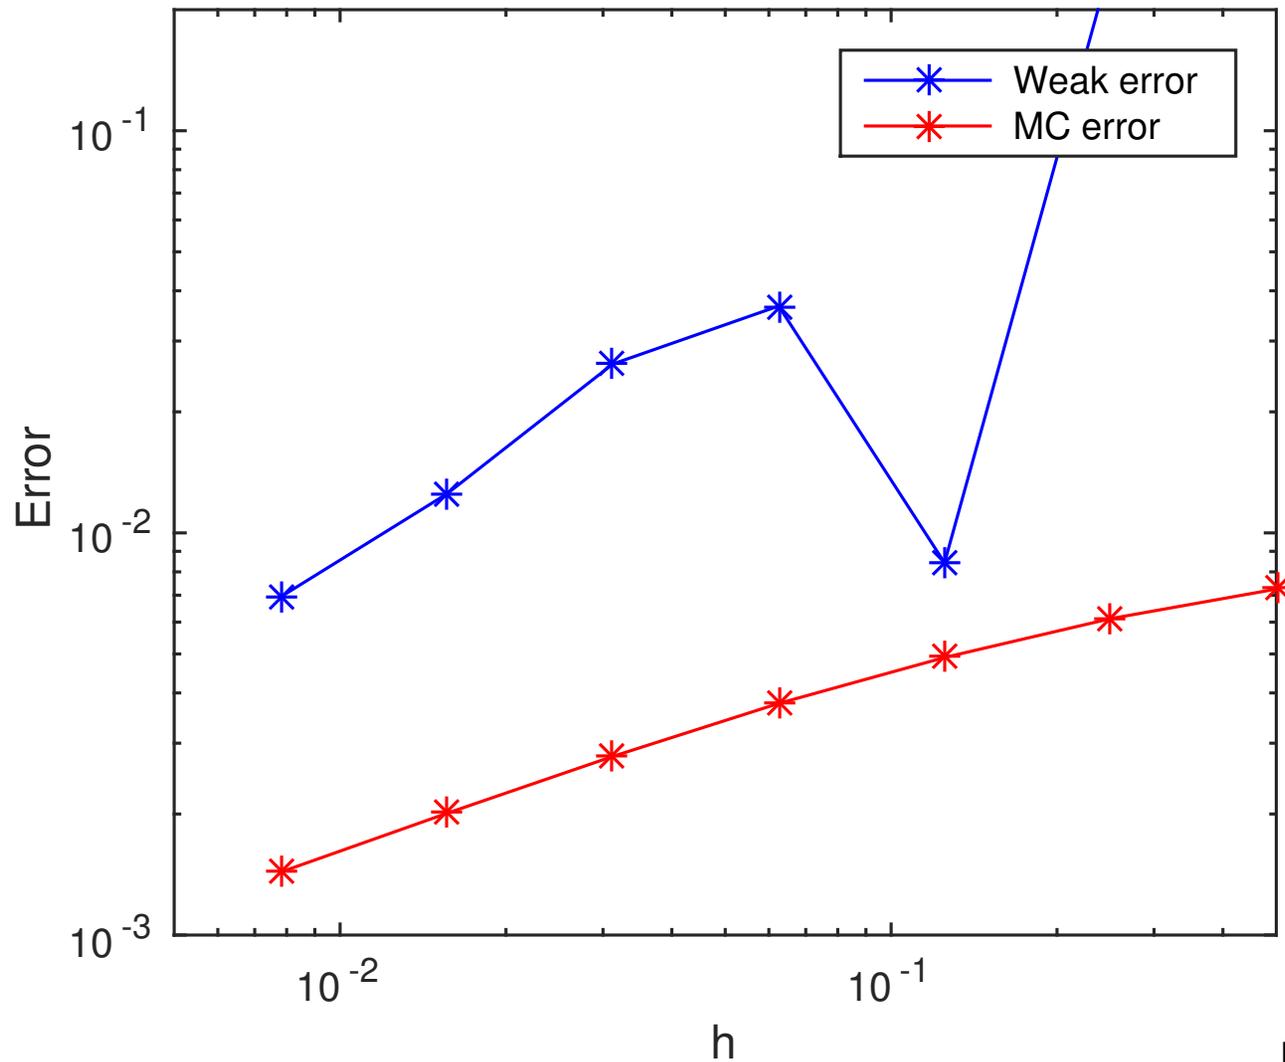
Weak convergence

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).

Weak convergence

Comparison to $2h$ approximation:



Barrier option

Some path-dependent options give only $O(\sqrt{h})$ weak convergence if the numerical payoff is not constructed carefully.

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

$$f = \exp(-rT) (\hat{S}_M - K)^+ \mathbf{1}_{\min_n \hat{S}_n > B}$$

Barrier option

Numerical demonstration: Geometric Brownian Motion

$$dS_t = r S_t dt + \sigma S_t dW_t$$

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

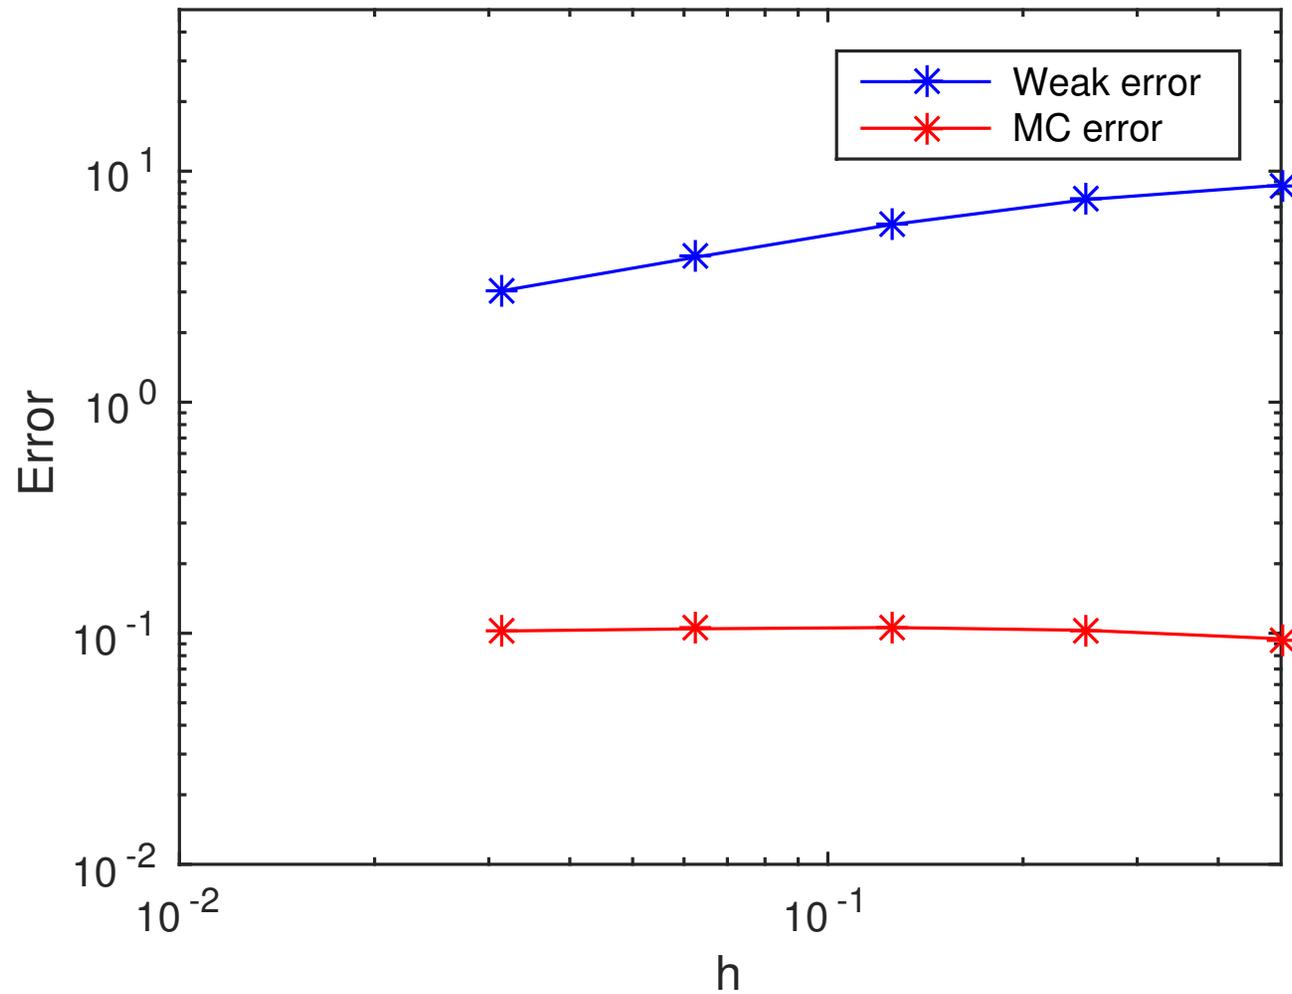
Down-and-out call: $S_0 = 100, K = 110, B = 90$.

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.)

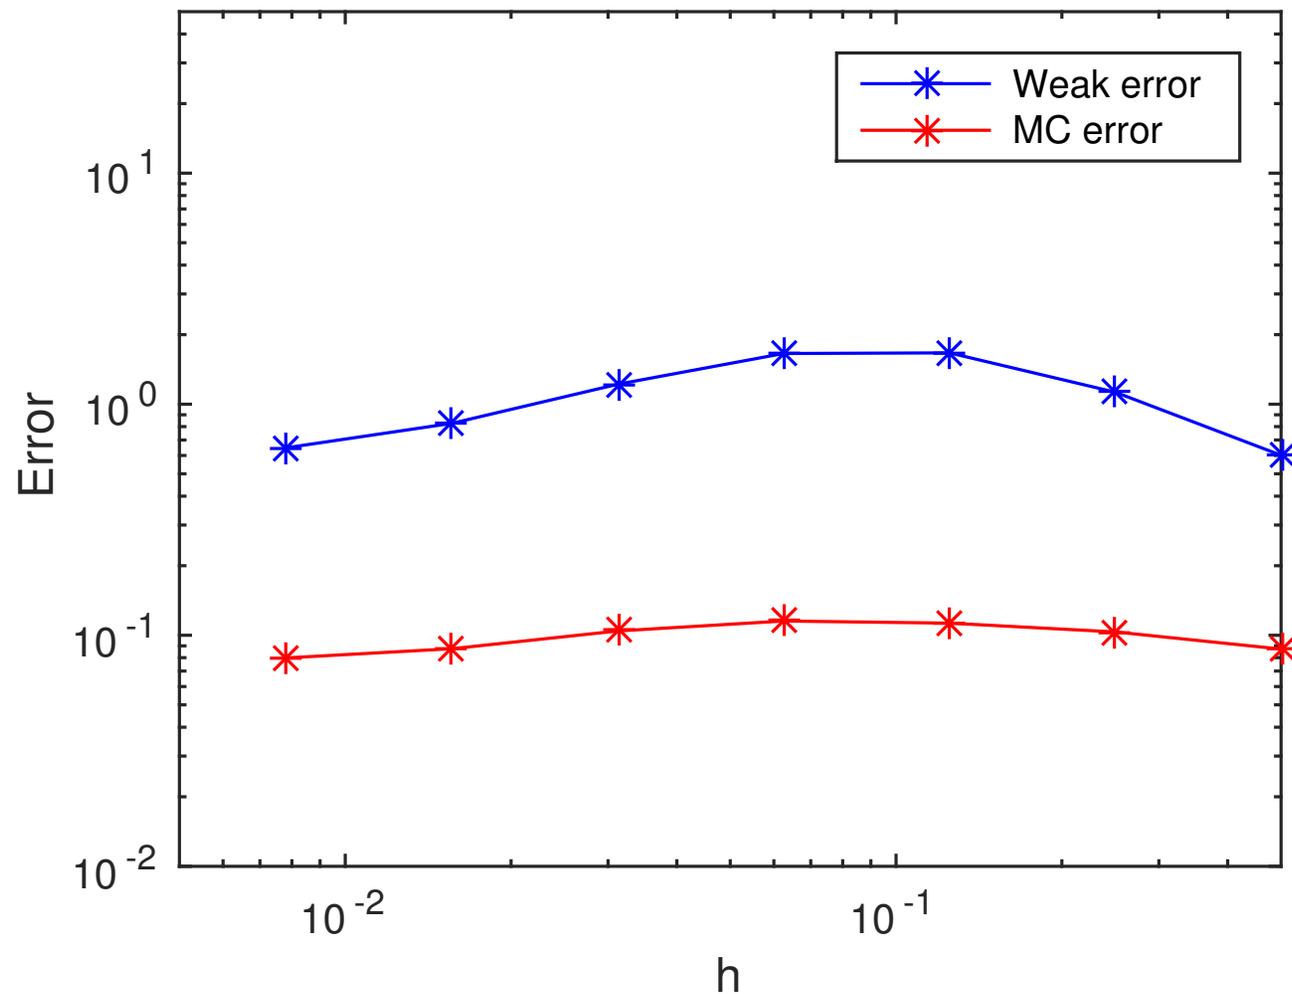
Barrier option

Comparison to exact solution:



Barrier option

Comparison to $2h$ approximation:



Lookback option

A floating-strike lookback call option has discounted payoff

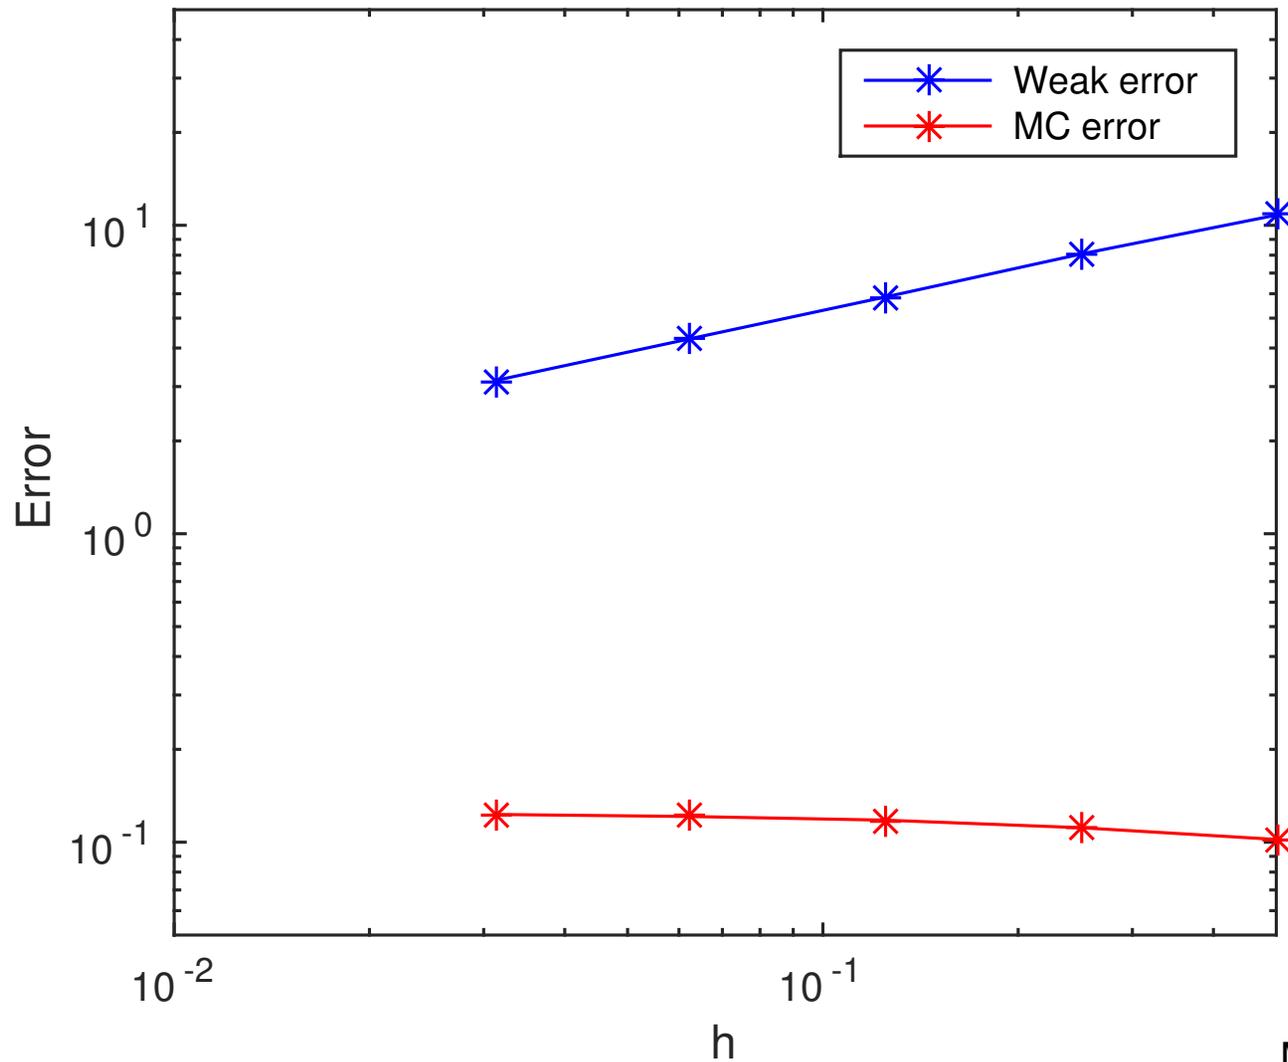
$$\exp(-rT) \left(S_T - \min_{[0,T]} S_t \right)$$

The natural numerical discretisation of this is

$$f = \exp(-rT) \left(\hat{S}_M - \min_n \hat{S}_n \right)$$

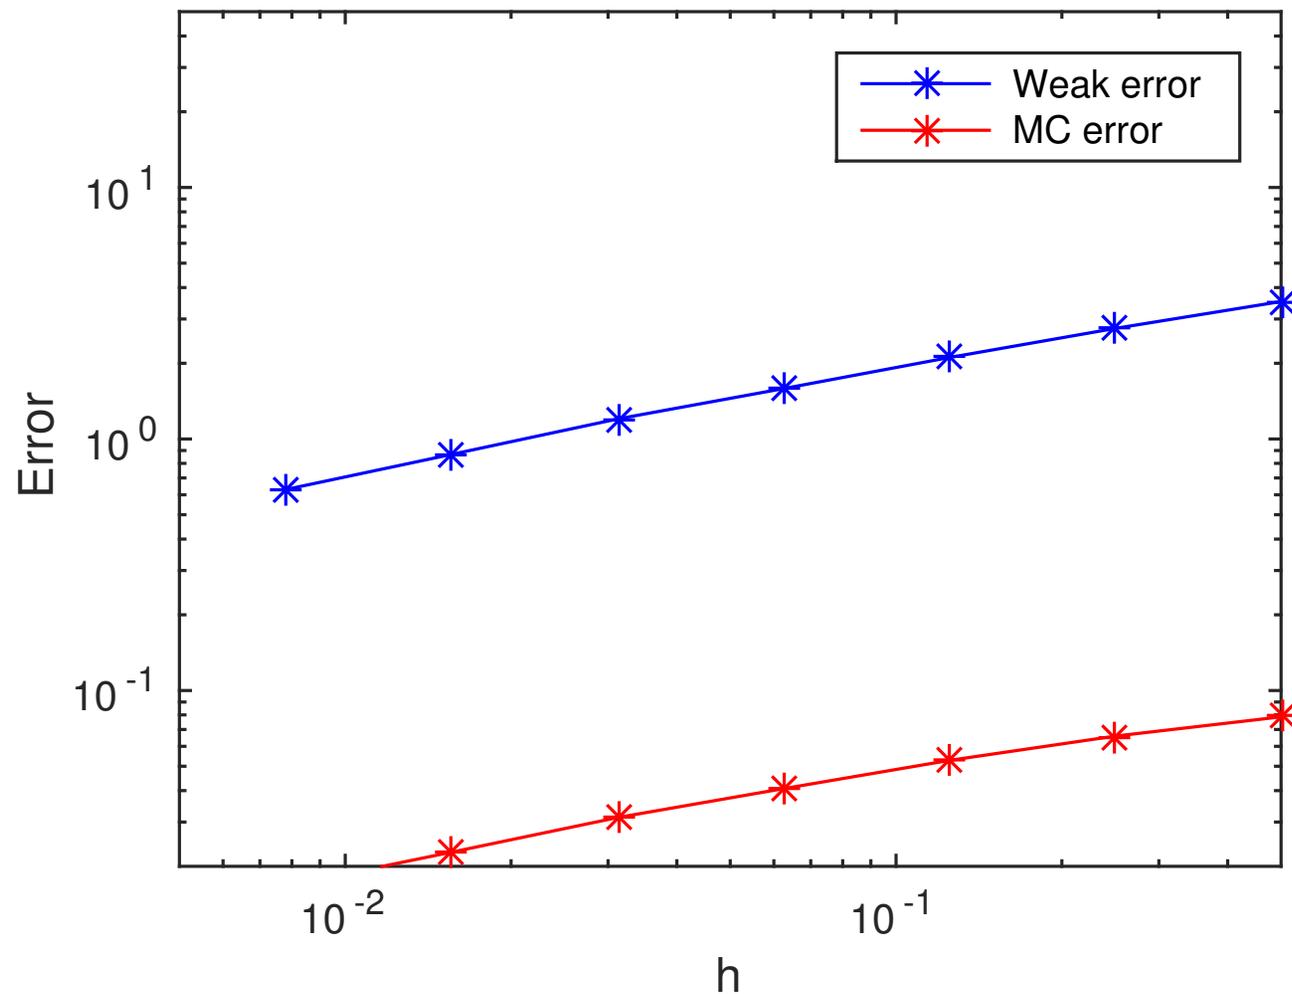
Lookback option

Comparison to exact solution



Lookback option

Comparison to $2h$ approximation



Brownian Bridge

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

Consider simple Brownian motion

$$dS = a dt + b dW$$

with constant a, b .

If we know the values S_n, S_{n+1} at times $t_n, t_{n+1} = t_n + h$, the Brownian Bridge construction considers the conditional behaviour of $S(t)$ within the timestep $t_n < t < t_{n+1}$.

Brownian Bridge

Two key results:

- The probability of the minimum dropping below B is:

$$\begin{aligned} \mathbb{P}\left(\min_{t_n < t < t_{n+1}} S(t) < B \mid S_n, S_{n+1} > B\right) \\ = \exp\left(-\frac{2(S_{n+1} - B)(S_n - B)}{b^2 h}\right) \end{aligned}$$

- A sample of the conditional minimum is given by

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

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

Barrier option

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

First, calculate \hat{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

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

Probability at end of not having crossed barrier is

$\prod_n (1 - P_n)$ and so the payoff is

$$\hat{f}(\hat{S}) = \exp(-rT) (\hat{S}_M - K)^+ \prod_n (1 - P_n).$$

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

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

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

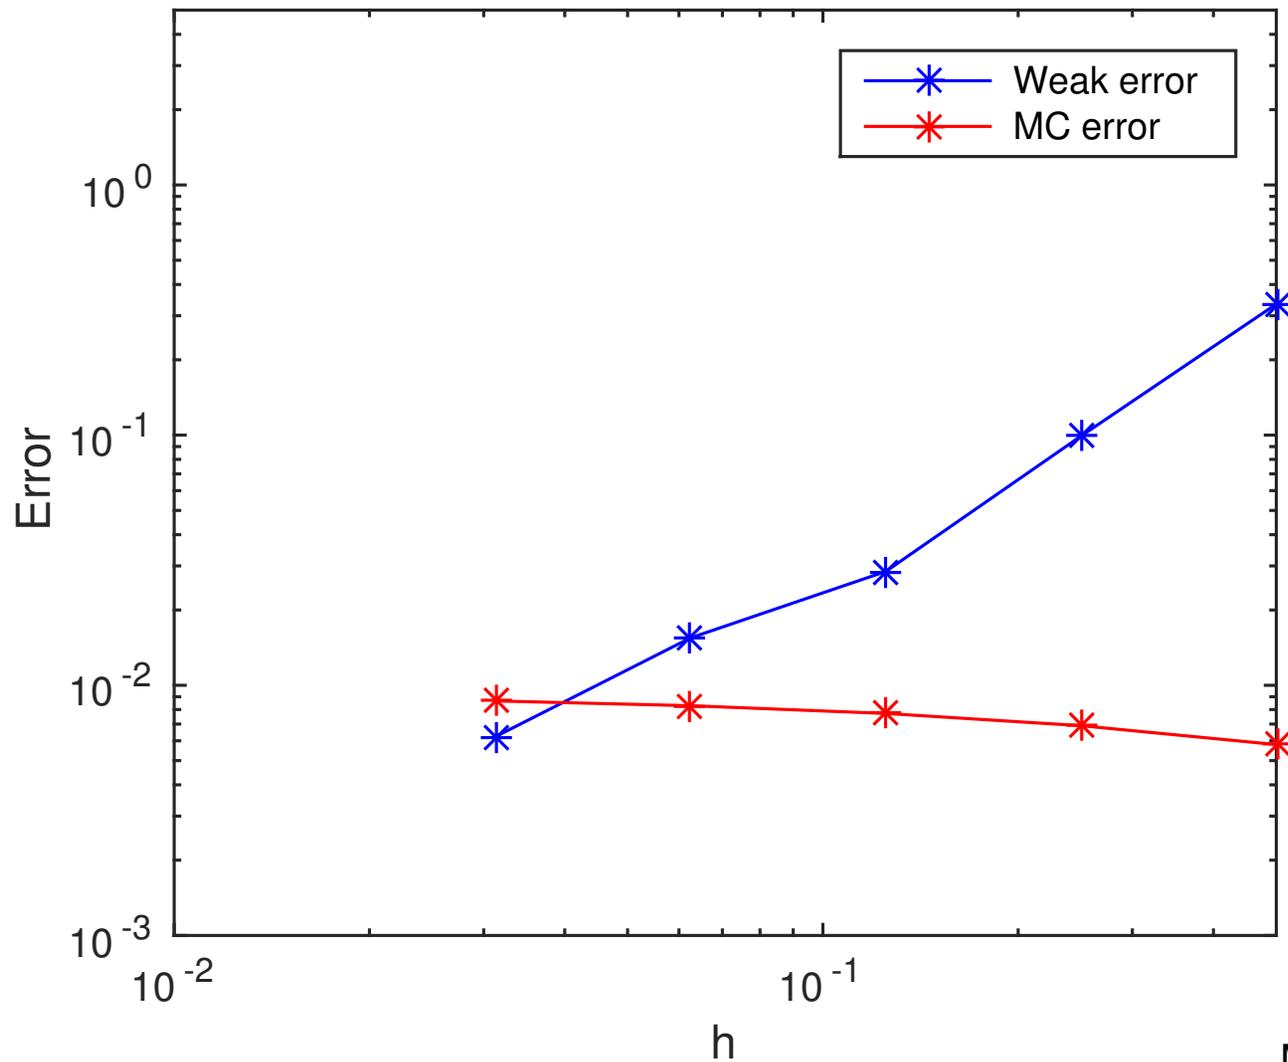
The payoff is then

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

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

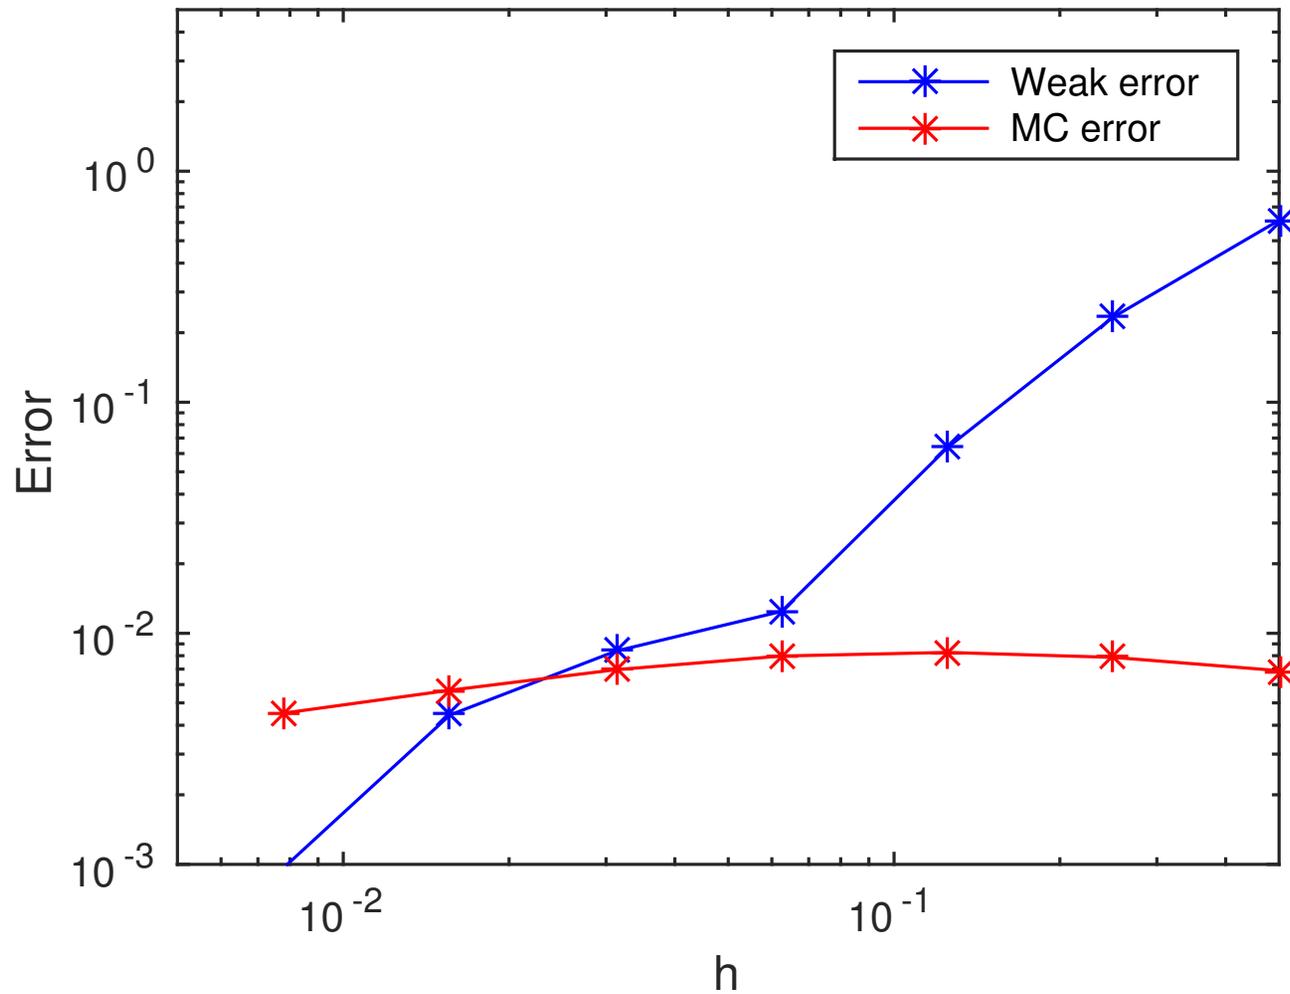
Weak convergence

Barrier: comparison to solution



Weak convergence

Barrier: h versus $2h$ solution



Lookback option

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

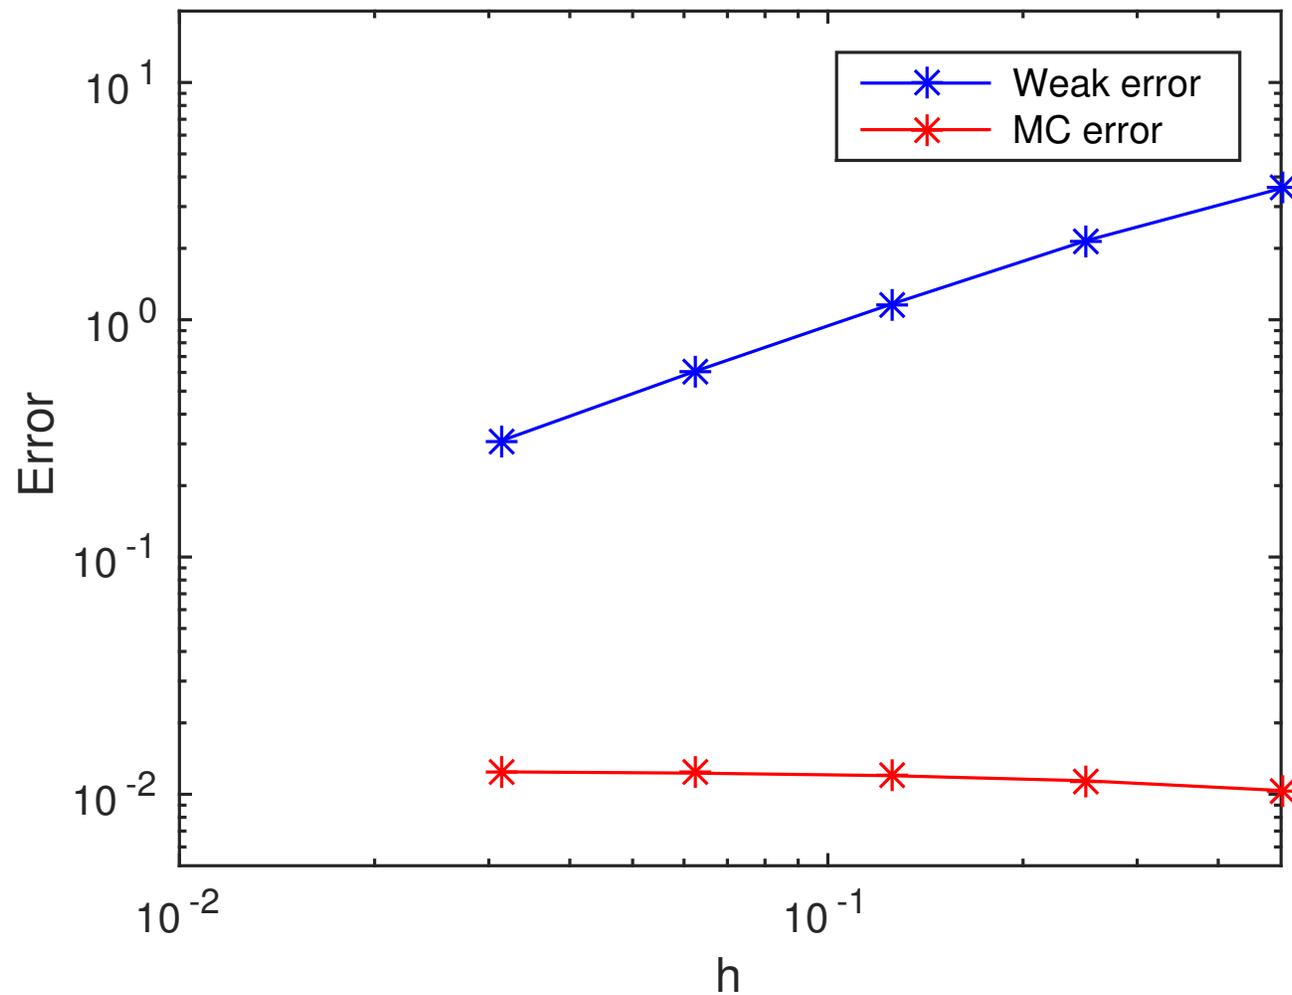
We construct a minimum \widehat{M}_n within each timestep and then the payoff is

$$\widehat{f}(\widehat{S}) = \exp(-rT) \left(\widehat{S}_M - \min_n \widehat{M}_n \right)$$

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

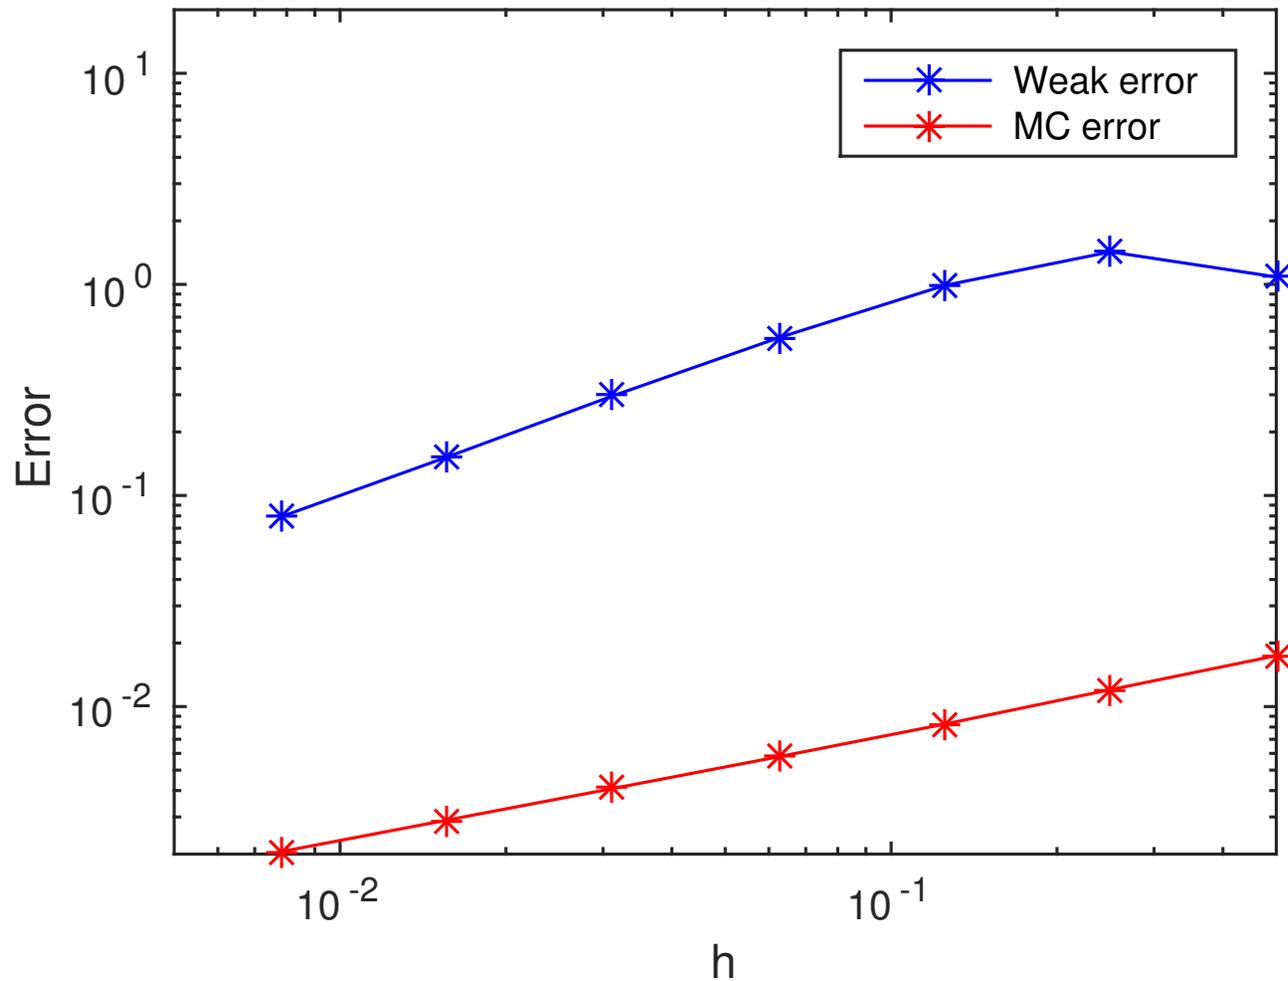
Weak convergence

Lookback: comparison to true solution



Weak convergence

Lookback: h versus $2h$ solution



Final Words

- The Greeks are very important
- “Bumping” is the standard approach, but Pathwise Sensitivity analysis is more efficient and more accurate, provided the payoff does not have a discontinuity
- The Euler-Maruyama method is used for path-dependent payoffs and general SDEs
- It gives $O(h)$ weak convergence – the error (bias) in the expected value
- Care must be taken with some path-dependent options to achieve this order of convergence.