

# Smoking adjoints, part II: fast Monte Carlo Greeks

Mike Giles

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

Oxford University Mathematical Institute  
Oxford-Man Institute of Quantitative Finance

Quant Congress '08

# “Smoking Adjoints”

Paper with Paul Glasserman in *Risk* in 2006 on the use of adjoints in computing pathwise sensitivities attracted a lot of interest, and questions:

- what is involved in practice in creating an adjoint code, and can it be simplified?
- what about barriers and American options?
- do we really have to differentiate the payoff?
- what about non-differentiable payoffs?

# Outline

- standard and adjoint pathwise sensitivities
- use of automatic differentiation ideas/tools
- barriers and American options
- “vibrato” Monte Carlo for non-differentiable payoffs

# Generic Problem

Stochastic differential equation with general drift and volatility terms:

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

For a simple European option we want to compute the expected discounted payoff value dependent on the terminal state:

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

Note: the drift and volatility functions are almost always differentiable, but the payoff  $f(S)$  is often not.

# Numerical discretisation

Euler discretisation with timestep  $h$ :

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

gives approximate expectation

$$\widehat{V} = \mathbb{E} \left[ f(\widehat{S}_N) \right]$$

# Pathwise sensitivity

Differentiating with respect to an arbitrary parameter  $\theta$  gives

$$\frac{\partial \hat{S}_{n+1}}{\partial \theta} = \left( 1 + \frac{\partial a}{\partial S} h + \frac{\partial b}{\partial S} \Delta W_n \right) \frac{\partial \hat{S}_n}{\partial \theta} + \frac{\partial a}{\partial \theta} h + \frac{\partial b}{\partial \theta} \Delta W_n$$

leading to

$$\frac{\partial \hat{V}}{\partial \theta} = \mathbb{E} \left[ \frac{\partial f}{\partial S}(\hat{S}_N) \frac{\partial \hat{S}_N}{\partial \theta} \right]$$

This is valid if  $a(S, \theta)$  and  $b(S, \theta)$  are differentiable, and the payoff  $f(S)$  is continuous and piecewise differentiable.

# Adjoint sensitivity

The adjoint approach is an efficient implementation of pathwise sensitivities – it gives exactly the same value.

Consider a process in which a vector input  $\alpha$  leads to a final state vector  $S$  which is used to compute a scalar payoff  $P$

$$\alpha \longrightarrow S \longrightarrow P$$

Taking  $\dot{\alpha}, \dot{S}, \dot{P}$  to be the derivatives w.r.t.  $j^{\text{th}}$  component of  $\alpha$ , then

$$\dot{S} = \frac{\partial S}{\partial \alpha} \dot{\alpha}, \quad \dot{P} = \frac{\partial P}{\partial S} \dot{S},$$

and hence

$$\dot{P} = \frac{\partial P}{\partial S} \frac{\partial S}{\partial \alpha} \dot{\alpha}.$$

# Adjoint sensitivity

Alternatively, defining  $\bar{\alpha}, \bar{S}, \bar{P}$  to be the derivatives of  $P$  with respect to  $\alpha, S, P$ , then

$$\bar{\alpha} \stackrel{\text{def}}{=} \left( \frac{\partial P}{\partial \alpha} \right)^T = \left( \frac{\partial P}{\partial S} \frac{\partial S}{\partial \alpha} \right)^T = \left( \frac{\partial S}{\partial \alpha} \right)^T \bar{S},$$

and similarly

$$\bar{S} = \left( \frac{\partial P}{\partial S} \right)^T \bar{P},$$

giving

$$\bar{\alpha} = \left( \frac{\partial S}{\partial \alpha} \right)^T \left( \frac{\partial P}{\partial S} \right)^T \bar{P}.$$



# Adjoint sensitivity

The two are mathematically equivalent, since

$$\dot{P} = \frac{\partial P}{\partial \alpha} \dot{\alpha} = \bar{\alpha}^T \dot{\alpha} = \bar{\alpha}_j$$

but the adjoint approach is much cheaper because a single calculation gives  $\bar{\alpha}$ , the sensitivity of  $P$  to each one of the elements of  $\alpha$ .

- standard approach: cost proportional to the number of Greeks
- adjoint approach: cost independent
- crossover point: 4 – 6 Greeks?

# Adjoint sensitivity

Note that the standard approach goes forward

$$\dot{\alpha} \longrightarrow \dot{S} \longrightarrow \dot{P}$$

while the adjoint approach does the reverse

$$\bar{\alpha} \longleftarrow \bar{S} \longleftarrow \bar{P}.$$

These correspond to the forward and reverse modes of Automatic Differentiation (AD).

“Smoking Adjoints” paper extended this to multiple timesteps in the path calculation — instead, we’ll extend it to the steps in a whole computer program.

# Automatic Differentiation

A computer instruction creates an additional new value:

$$\mathbf{u}^n = \mathbf{f}^n(\mathbf{u}^{n-1}) \equiv \begin{pmatrix} \mathbf{u}^{n-1} \\ f_n(\mathbf{u}^{n-1}) \end{pmatrix},$$

A computer program is the composition of  $N$  such steps:

$$\mathbf{u}^N = \mathbf{f}^N \circ \mathbf{f}^{N-1} \circ \dots \circ \mathbf{f}^2 \circ \mathbf{f}^1(\mathbf{u}^0)$$

Differentiation w.r.t. one element of the input vector gives

$$\dot{\mathbf{u}}^N = D^N D^{N-1} \dots D^2 D^1 \dot{\mathbf{u}}^0, \quad D^n \equiv \begin{pmatrix} I^{n-1} \\ \partial f_n / \partial \mathbf{u}^{n-1} \end{pmatrix}$$

# Automatic Differentiation

In reverse mode, we consider the sensitivity of one element of the output vector, to get

$$\begin{aligned}(\bar{\mathbf{u}}^{n-1})^T &\equiv \frac{\partial u_i^N}{\partial \mathbf{u}^{n-1}} = \frac{\partial u_i^N}{\partial \mathbf{u}^n} \frac{\partial \mathbf{u}^n}{\partial \mathbf{u}^{n-1}} = (\bar{\mathbf{u}}^n)^T D^n, \\ &\implies \bar{\mathbf{u}}^{n-1} = (D^n)^T \bar{\mathbf{u}}^n.\end{aligned}$$

and hence

$$\bar{\mathbf{u}}^0 = (D^1)^T (D^2)^T \dots (D^{N-1})^T (D^N)^T \bar{\mathbf{u}}^N.$$

Note: need to go forward through original calculation to compute/store the  $D^n$ , then go in reverse to compute  $\bar{\mathbf{u}}^n$

# Automatic Differentiation

This gives a prescriptive algorithm for reverse mode differentiation.

Again the reverse mode is much more efficient if we want the sensitivity of a single output to multiple inputs.

Key result is that the cost of the reverse mode is at worst a factor 4 greater than the cost of the original calculation, regardless of how many sensitivities are being computed!

The storage of the  $D^n$  is minor for SDEs – much more of a concern for PDEs.

# Automatic Differentiation

Manual implementation of the forward/reverse mode algorithms is possible but a little tedious.

Fortunately, automated tools have been developed, following one of two approaches:


- operator overloading (ADOL-C, FADBAD++)
  - well developed, robust, not very efficient
- source code transformation (Tapenade, TAC)
  - still under development for C, major challenges posed by C++, often close to efficiency of hand-coded adjoints

# LIBOR Application

- testcase from “Smoking Adjoints” paper
- test problem performs  $N$  timesteps with a vector of  $N + 40$  forward rates, and computes the  $N + 40$  deltas and vegas for a portfolio of swaptions
- hand-coded using the ideas from automatic differentiation

# LIBOR Application


Finite differences versus forward pathwise sensitivities:





# LIBOR Application

Hand-coded forward versus adjoint pathwise sensitivities:



# Barrier options

The big limitation of pathwise sensitivity calculations (and their adjoint equivalent) is that the validity depends on the payoff being continuous and piecewise differentiable.

What about a barrier option?

A simple implementation uses the Euler discretisation and defines the numerical payoff of a down-and-out call as

$$\hat{P} = e^{-rT} (\hat{S}_N - K)^+ \mathbf{1}(\min_n \hat{S}_n > B)$$

This is discontinuous when the smallest  $\hat{S}_n$  crosses  $B$ .

# Barrier options

It is better to define the payoff as

$$\hat{P} = e^{-rT} (\hat{S}_N - K)^+ \prod_n (1 - p_n)$$

where

$$p_n = \exp \left( \frac{-2 (\hat{S}_n - B)^+ (\hat{S}_{n+1} - B)^+}{b_n^2 h} \right)$$

is the approximate probability that the path crosses the barrier in time interval  $[t_n, t_{n+1}]$ , conditional on  $\hat{S}_n, \hat{S}_{n+1}$ .

This is more accurate (bias is  $O(h)$  instead of  $O(h^{1/2})$ ), and continuous as  $\hat{S}_n$  crosses the barrier.

# American options

Longstaff-Schwartz approach:

- calculate a set of paths to maturity
- step backwards in time from maturity
  - calculate least-squares approximation of continuation value
  - if best to exercise, set value equal to exercise value
  - otherwise, discount existing path value

# American options

- small discontinuity because of difference between continuation value and discounted path value
- suspect the error due to this is small (negligible?) but haven't yet investigated
- pathwise sensitivity analysis is equivalent to fixing the exercise boundary, which has zero effect in PDE formulation
- could be viewed as naive barrier option treatment; use probabilistic treatment to regain continuity
- alternatively, use Tsitsiklis & Van Roy treatment; set path value to approximate continuation value

# Vibrato Monte Carlo

One remaining problem – what if payoff is not differentiable?

- Likelihood Ratio Method (LRM)
  - estimator variance proportional to  $h^{-1}$
- Malliavin calculus
  - recent paper by Glasserman and Chen shows it can be viewed as a pathwise/LRM hybrid
  - might be good choice when few Greeks needed
- new “vibrato” Monte Carlo idea
  - also a pathwise/LRM hybrid
  - variance proportional to  $h^{-1/2}$
  - efficient adjoint implementation

# Vibrato Monte Carlo

- new idea, based on conditional expectation for a simple digital option (Glasserman's MC book)
- output of each SDE path calculation becomes a narrow (multivariate) Normal distribution
- combine pathwise sensitivity for the differentiable SDE, with LRM for the non-differentiable payoff
- avoiding the differentiation of the payoff also simplifies the implementation in real-world setting

# Vibrato Monte Carlo

Final timestep of Euler path discretisation is

$$\hat{S}_N = \hat{S}_{N-1} + a(\hat{S}_{N-1}, t_{N-1}) h + b(\hat{S}_{N-1}, t_{N-1}) \Delta W_{N-1}$$

Instead of using random number generator to get a value for  $\Delta W_{N-1}$ , consider the whole distribution of possible values, so  $\hat{S}_N$  has a Normal distribution with mean

$$\mu(W) = \hat{S}_{N-1} + a(\hat{S}_{N-1}, t_{N-1}) h$$

and standard deviation

$$\sigma(W) = b(\hat{S}_{N-1}, t_{N-1}) \sqrt{h}$$

where  $W \equiv (\Delta W_0, \Delta W_1, \dots, \Delta W_{N-2})$ .



# Vibrato Monte Carlo

For a particular path given by a particular vector  $W$ , the expected payoff is

$$\mathbb{E}_Z[f(\mu + \sigma Z)]$$

where  $Z$  is a Normal random variable with zero mean and unit variance.

Averaging over all  $W$  then gives the same overall expectation as before.

Note also that, for given  $W$ ,  $\hat{S}_N$  has a Normal distribution with

$$p_S(\hat{S}_N) = \frac{1}{\sqrt{2\pi} \sigma} \exp\left(-\frac{(\hat{S}_N - \mu)^2}{2\sigma^2}\right)$$

# Vibrato Monte Carlo

In the case of a simple digital call with strike  $K$ , can use the analytic solution

$$\mathbb{E}_Z[f(\mu + \sigma Z)] = \exp(-rT) \Phi\left(\frac{\mu - K}{\sigma}\right).$$

- for each  $W$ , the payoff is now smooth, differentiable
- derivative is  $O(h^{-1/2})$  near strike, near zero elsewhere  
 $\implies$  variance is  $O(h^{-1/2})$

Analytic evaluation of conditional expectation not possible in general for multivariate cases, so use Monte Carlo estimation!

# Vibrato Monte Carlo

Main novelty comes in calculating the sensitivity.

For a particular  $W$ , we have a Normal probability distribution for  $\hat{S}_N$  and can apply the Likelihood Ratio method to get

$$\frac{\partial}{\partial \theta} \mathbb{E}_Z \left[ f(\hat{S}_N) \right] = \mathbb{E}_Z \left[ f(\hat{S}_N) \frac{\partial(\log p_S)}{\partial \theta} \right],$$

where

$$\begin{aligned} \frac{\partial(\log p_S)}{\partial \theta} &= \frac{\partial(\log p_S)}{\partial \mu} \frac{\partial \mu}{\partial \theta} + \frac{\partial(\log p_S)}{\partial \sigma} \frac{\partial \sigma}{\partial \theta} \\ &= \frac{Z}{\sigma} \frac{\partial \mu}{\partial \theta} + \frac{Z^2 - 1}{\sigma} \frac{\partial \sigma}{\partial \theta}. \end{aligned}$$

Averaging over all  $W$  then gives the expected sensitivity.

# Vibrato Monte Carlo

To improve the variance, we note that

$$\mathbb{E}[1] = 1 \quad \implies \quad \mathbb{E}_Z \left[ \frac{\partial(\log p_S)}{\partial \theta} \right] = 0$$

and hence

$$\frac{\partial}{\partial \theta} \mathbb{E}_Z \left[ f(\hat{S}_N) \right] = \mathbb{E}_Z \left[ \left( f(\mu + \sigma Z) - f(\mu) \right) \frac{\partial(\log p_S)}{\partial \theta} \right].$$

The quantity

$$\hat{P} = \left( f(\mu + \sigma Z) - f(\mu) \right) \frac{\partial(\log p_S)}{\partial \theta}$$

has  $O(1)$  variance when  $f(S)$  is Lipschitz.

# Vibrato Monte Carlo

In the multivariate extension,

$$\hat{S}(W, Z) = \mu + C Z$$

where  $\mu$  is the mean,  $\Sigma = C C^T$  is the variance, and  $Z$  is a vector of uncorrelated Normals. The joint p.d.f. is

$$\log p_S = -\frac{1}{2} \log |\Sigma| - \frac{1}{2} (\hat{S} - \mu)^T \Sigma^{-1} (\hat{S} - \mu) - \frac{1}{2} d \log(2\pi).$$

$$\implies \frac{\partial \log p_S}{\partial \mu} = C^{-T} Z, \quad \frac{\partial \log p_S}{\partial \Sigma} = \frac{1}{2} C^{-T} (Z Z^T - I) C^{-1}$$

Can also handle options dependent on values at intermediate times by using Brownian interpolation between simulation times on either side.

# Vibrato Monte Carlo

For each  $W$ , in forward mode we have

$$\alpha, \dot{\alpha} \longrightarrow \mu, \dot{\mu}, \Sigma, \dot{\Sigma} \longrightarrow \text{payoff + sensitivity}$$

- first bit – pathwise sensitivity calculation
- second bit – Likelihood Ratio Method

For maximum efficiency can use adjoint/reverse mode

$$\begin{array}{ccccccc} \alpha & \longrightarrow & \mu, \Sigma & \longrightarrow & \text{payoff} \\ \bar{\alpha} & \longleftarrow & \bar{\mu}, \bar{\Sigma} & \longleftarrow & \text{sensitivity} \end{array}$$

$\bar{\mu}, \bar{\Sigma}$  are coefficients multiplying  $\dot{\mu}, \dot{\Sigma}$  in forward mode

# Vibrato Monte Carlo

Test case: Geometric Brownian motion

$$dS_t^{(1)} = r S_t^{(1)} dt + \sigma^{(1)} S_t^{(1)} dW_t^{(1)}$$

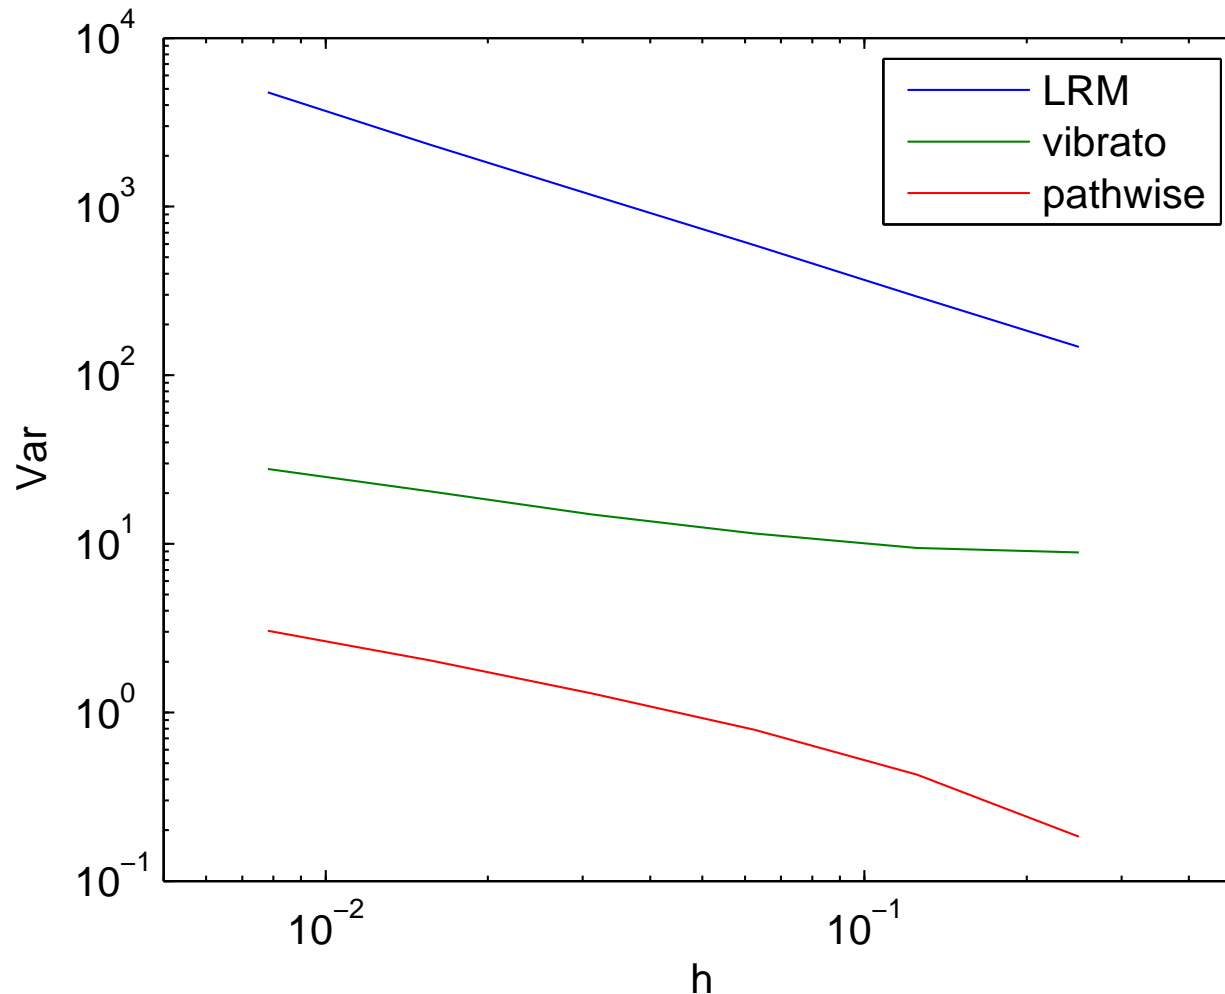
$$dS_t^{(2)} = r S_t^{(2)} dt + \sigma^{(2)} S_t^{(2)} dW_t^{(2)}$$

with a simple digital call option based solely on  $S_T^{(1)}$ .

Parameters:  $r = 0.05$ ,  $\sigma^{(1)} = 0.2$ ,  $\sigma^{(2)} = 0.3$ ,  $T = 1$ ,  
 $S_0^{(1)} = S_0^{(2)} = 100$ ,  $K = 100$ ,  $\rho = 0.5$

Numerical results compare LRM, vibrato with one  $Z$  per  $W$ , and pathwise using conditional expectation.

# Vibrato Monte Carlo



More  $Z$  samples per path would bring the vibrato results closer to the pathwise results based on analytic expectation



# Conclusions

- Adjoint implementation of pathwise sensitivities is very efficient when multiple Greeks are needed
- Automatic differentiation tools can aid software development
- Barrier and American options can be handled with care
- New “vibrato” MC idea can handle discontinuous payoffs and avoid the need to differentiate payoffs

Further information:

- `people.maths.ox.ac.uk/~gilesm/`
- Email: `mike.giles@maths.ox.ac.uk`