

Smoking adjoints: fast Monte Carlo Greeks

Monte Carlo calculation of price sensitivities for hedging is often very time-consuming. Michael Giles and Paul Glasserman develop an adjoint method to accelerate the calculation. The method is particularly effective in estimating sensitivities to a large number of inputs, such as initial rates on a forward curve or points on a volatility surface. The authors apply the method to the Libor market model and show that it is much faster than previous methods

The efficient calculation of price sensitivities continues to be among the greatest practical challenges facing users of Monte Carlo methods in the derivatives industry. Computing Greeks is essential to hedging and risk management, but typically requires substantially more computing time than pricing a derivative. This article shows how an adjoint formulation can be used to accelerate the calculation of the Greeks. This method is particularly well suited to applications requiring sensitivities to a large number of parameters. Examples include interest rate derivatives requiring sensitivities to all initial forward rates and equity derivatives requiring sensitivities to all points on a volatility surface.

The simplest methods for estimating Greeks are based on finite difference approximations, in which a Monte Carlo pricing routine is rerun multiple times at different settings of the input parameters in order to estimate sensitivities to the parameters. In the fixed-income setting, for example, this would mean perturbing each initial forward rate and then rerunning the Monte Carlo simulation to re-price a security or a whole book. The main virtues of this method are that it is straightforward to understand and requires no additional programming. But the bias and variance properties of finite difference estimates can be rather poor, and their computing time requirements grow with the number of input parameters.

Better estimates of price sensitivities can often be derived by using information about model dynamics in a Monte Carlo simulation. Techniques for doing this include the pathwise method and likelihood ratio method, both of which are reviewed in chapter 7 of Glasserman (2004). When applicable, these methods produce unbiased estimates of price sensitivities from a single set of simulated paths, that is, without perturbing any parameters. The pathwise method accomplishes this by differentiating the evolution of the underlying assets or state variables along each path; the likelihood ratio method instead differentiates the transition density of the underlying assets or state variables. In comparison with finite difference estimates, these methods require additional model analysis and programming, but the additional effort is often justified by the improvement in the quality of calculated Greeks.

The adjoint method we develop here applies ideas used in computational fluid dynamics (Giles & Pierce, 2000) to the calculation of pathwise estimates of Greeks. The estimate calculated using the adjoint method is identical to the ordinary pathwise estimate; its potential advantage is therefore computational, rather than statistical. The relative merits of the ordinary (forward) calculation of pathwise Greeks and the adjoint calculation can be summarised as follows: a) the adjoint method is advantageous for calculating the sensitivities of a small number of securities with respect to a large number of parameters; and b) the forward method is advantageous for calculating the sensitivities of many securities with respect to a small

number of parameters. The 'small number of securities' in this dichotomy could be an entire book, consisting of many individual securities, so long as the sensitivities to be calculated are for the book as a whole and not for the constituent securities.

The rest of this article is organised as follows. The next section reviews the usual forward calculation of pathwise Greeks and the subsequent section illustrates its application in the Libor market model. We then develop the adjoint method for delta estimates, and extend it to applications such as vega estimation requiring sensitivities to parameters of model dynamics, rather than just sensitivities to initial conditions. We then extend it to gamma estimation. We use the Libor market model as an illustrative example in both settings. Lastly, we present numerical results that illustrate the computational savings offered by the adjoint method.

Pathwise delta: forward method

We start by reviewing the application of the pathwise method for computing price sensitivities in the setting of a multi-dimensional diffusion process satisfying a stochastic differential equation:

$$d\tilde{X}(t) = a(\tilde{X}(t))dt + b(\tilde{X}(t))dW(t) \quad (1)$$

The process \tilde{X} is m -dimensional, W is a d -dimensional Brownian motion, $a(\cdot)$ takes values in R^m and $b(\cdot)$ takes values in $R^{m \times d}$. For example, \tilde{X} could record a vector of equity prices or – as in the case of the Libor market model, below – a vector of forward rates. We take (1) to be the risk-neutral or otherwise risk-adjusted dynamics of the relevant financial variables. A derivative maturing at time T with discounted payout $g(\tilde{X}(T))$ has price $E[g(\tilde{X}(T))]$, the expected value of the discounted payout.

In a Monte Carlo simulation, the evolution of the process \tilde{X} is usually approximated using an Euler scheme. For simplicity, we take a fixed time step $h = T/N$, with N an integer. We write $X(n)$ for the Euler approximation at time nh , which evolves according to:

$$X(n+1) = X(n) + a(X(n))h + b(X(n))Z(n+1)\sqrt{h}, \quad X(0) = \tilde{X}(0) \quad (2)$$


where $Z(1), Z(2), \dots$ are independent d -dimensional standard normal random vectors. With the normal random variables held fixed, (2) takes the form:

$$X(n+1) = F_n(X(n)) \quad (3)$$

with F_n a transformation from R^m to R^m .

The price of the derivative with discounted payout function g is estimated using the average of independent replications of $g(X(N))$, $N = T/h$.

2. Data flow showing relationship between forward and adjoint calculations



$$\Delta_{ij}(n) = \frac{\partial L_i(n)}{\partial L_j(0)}, \quad i = 1, \dots, m, \quad j = 1, \dots, i$$

which can be found by differentiating (7). In the notation of (5), the matrix $D(n)$ has the structure shown in figure 1, with diagonal entries:

$$D_{ii}(n) = \begin{cases} 1 & i < \eta(nh) \\ \frac{L_i(n+1)}{L_i(n)} + \frac{L_i(n+1) \|\sigma_i\|^2 \delta_i h}{(1 + \delta_i L_i(n))^2} & i \geq \eta(nh) \end{cases}$$

and, for $j \neq i$:

$$D_{ij}(n) = \begin{cases} \frac{L_i(n+1) \sigma_i^T \sigma_j \delta_j h}{(1 + \delta_j L_j(n))^2} & i > j \geq \eta(nh) \\ 0 & \text{otherwise} \end{cases}$$

The efficient implementation used in the numerical results of Glasserman & Zhao (1999) uses $\Delta_{ij}(n+1) = \Delta_{ij}(n)$ for $i < \eta(nh)$, while for $i \geq \eta(nh)$:

$$\Delta_{ij}(n+1) = \frac{L_i(n+1)}{L_i(n)} \Delta_{ij}(n) + L_i(n+1) \sigma_i^T \sum_{k=\eta(nh)}^i \frac{\sigma_k \delta_k h \Delta_{kj}(n)}{(1 + \delta_k L_k(n))^2}$$

The summations on the right can be computed at a cost that is $O(m)$ for each j , and hence the total computational cost per time step is $O(m^2)$ rather than the $O(m^3)$ cost of implementing (5) in general.

Despite this, the number of forward rates m in the Libor market model can easily be 20–80, making the numerical evaluation of $\Delta_{ij}(n)$ rather costly. To get around this problem, Glasserman & Zhao (1999) proposed faster approximations to (5). The adjoint method in the next section can achieve computational savings without introducing any approximation beyond that already present in the Euler scheme.

Pathwise delta: adjoint method

Consider again the general setting of (1) and (2) and write $\partial g / \partial X(0)$ for the row vector of derivatives of $g(X(N))$ with respect to the elements of $X(0)$. With (4) and (5), we can write this as:

$$\begin{aligned} \frac{\partial g}{\partial X(0)} &= \frac{\partial g}{\partial X(N)} \Delta(N) \\ &= \frac{\partial g}{\partial X(N)} D(N-1) D(N-2) \dots D(0) \Delta(0) \\ &= V(0)^T \Delta(0) \end{aligned} \quad (9)$$

where $V(0)$ can be calculated recursively using:

$$V(n) = D(n)^T V(n+1), \quad V(N) = \left(\frac{\partial g}{\partial X(N)} \right)^T \quad (10)$$

The key point is that the adjoint relation (10) is a vector recursion whereas (5) is a matrix recursion. Thus, rather than update m^2 variables at each time step, it suffices to update the m entries of the adjoint variables $V(n)$. This can represent a substantial saving.

The adjoint method accomplishes this by fixing the payout g in the initialisation of $V(N)$, whereas the forward method allows calculation of pathwise deltas for multiple payouts once the $\Delta(n)$ matrices have been simulated. Thus, the adjoint method is beneficial if we are interested in calculating sensitivities of a single function g with respect to multiple changes in the initial condition $X(0)$ – for example, if we need sensitivities with respect to each $X_i(0)$. The function g need not be associated with an individual security; it could be the value of an entire portfolio.

Equation (10) is a discrete adjoint to equation (5) in the same way that the differential equation $-dv/dt = A^T v$ is the adjoint counterpart to $du/dt = Au$, where u and v are in R^m and A is in $R^{m \times m}$. The terminology ‘discrete adjoint’ is used in computational engineering (Giles & Pierce, 2000), but in the computational finance context it might also be referred to as a ‘backward’ counterpart to the original forward pathwise sensitivity calculation since the adjoint recursion in (10) runs backward in time, starting at $V(N)$ and working recursively back to $V(0)$. To implement it, we need to store the vectors $X(0), \dots, X(N)$ as we simulate forward in time so that we can evaluate the matrices $D(N-1), \dots, D(0)$ as we work backward as illustrated in figure 2. This introduces some additional storage requirements, but these requirements are relatively minor because it suffices to store just the current path. The final calculation $V(0)^T \Delta(0)$ produces exactly the same result as the forward calculations (4)–(5), but it does so with $O(Nm^2)$ operations rather than $O(Nm^3)$ operations.

To help fix ideas, we unravel the adjoint calculation in the setting of the Libor market model. After initialising $V(N)$ according to (10), we set

Guidelines for the submission of technical articles

Risk welcomes the submission of technical articles on topics relevant to our readership. Core areas include market and credit risk measurement and management, the pricing and hedging of derivatives and/or structured securities, and the theoretical modelling and empirical observation of markets and portfolios. This list is not an exhaustive one.

The most important publication criteria are originality, exclusivity and relevance – we attempt to strike a balance between these. Given that *Risk* technical articles are shorter than those in dedicated academic journals, clarity of exposition is another yardstick for publication. Once received by the technical editor and his team, submissions are logged, and checked against the criteria above. Articles that fail to meet the criteria are rejected at this stage.

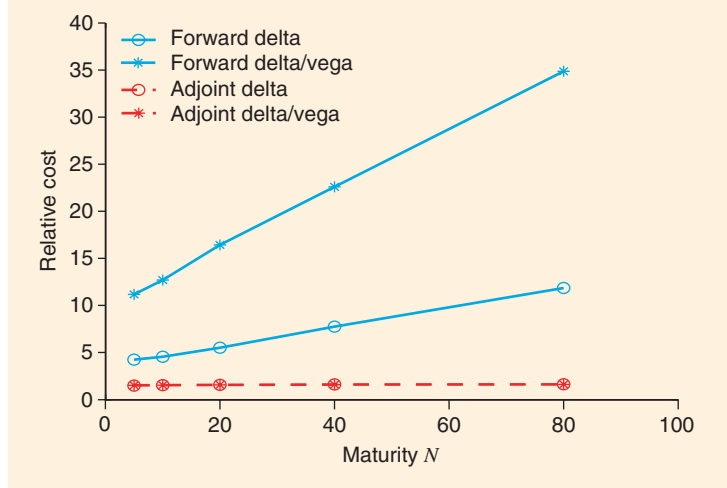
Articles are then sent to one or more anonymous referees for peer review. Our referees are drawn from the research groups, risk management departments and trading desks of major financial institutions, in addition to academia. Many have already published articles in *Risk*. Depending on the feedback from referees, the

technical editor makes a decision to reject or accept the submitted article. His decision is final.

We also welcome the submission of brief communications. These are also peer-reviewed contributions to *Risk* but the process is less formal than for full-length technical articles. Typically, brief communications address an extension or implementation issue arising from a full-length article that, while satisfying our originality, exclusivity and relevance requirements, does not deserve full-length treatment.

Submissions should be sent to the technical team at technical@incisivemedia.com. The preferred format is MS Word, although Adobe PDFs are acceptable. The maximum recommended length for articles is 3,500 words, and for brief communications 1,000 words, with some allowance for charts and/or formulas. We expect all articles and communications to contain references to previous literature. We reserve the right to cut accepted articles to satisfy production considerations. Authors should allow four to eight weeks for the refereeing process.

4. Relative CPU cost of forward and adjoint delta and vega evaluation for a portfolio of 15 swaptions



$$E_{i\ell m}(n) = \frac{\partial^2 F_i(n)}{\partial X_\ell(n) \partial X_m(n)}$$

For a particular index pair (j, k) , by defining:

$$G_i(n) = \Gamma_{ijk}(n), \quad C_i(n) = \sum_{\ell=1}^m \sum_{m=1}^m E_{i\ell m}(n) \Delta_{\ell j}(n) \Delta_{mk}(n)$$

this may be written as:

$$G(n+1) = D(n)G(n) + C(n)$$

This is now in exactly the same form as the vega calculation, and so the same adjoint approach can be used. Option payouts ordinarily fail to be twice differentiable, so using (13) requires replacing the true payout g with a smoothed approximation; this is the subject of current research.

The computational operation count is $O(Nm^3)$ for the forward calculation of $L(n)$ and $\Delta(n)$ (and hence $D(n)$ and the vectors $C(n)$ for each index pair (j, k)) plus $O(Nm^2)$ for the backward calculation of the adjoint variables $V(n)$, followed by an $O(Nm^3)$ cost for evaluating the final sums in (12) for each (j, k) . This is again a factor $O(m)$ less expensive than the alternative approach based on a forward calculation of $\Gamma_{ijk}(n)$.

Numerical results

Since the adjoint method produces exactly the same sensitivity values as the forward pathwise approach, the numerical results address the computational savings given by the adjoint approach applied to the Libor market model. The calculations are performed using one time step per Libor interval (that is, the time step h equals the spacing $\delta_i \equiv \delta$, which we take to be a quarter of a year). We take the initial forward curve to be flat at 5% and all volatilities equal to 20% in a single-factor ($d = 1$) model. Our test portfolio consists of options on one-year, two-year, five-year, seven-year and 10-year swaps with quarterly payments and swap rates of 4.5%, 5.0% and 5.5%, for a total of 15 swaptions. All swaptions expire in N periods, with N varying from one to 80.

Figure 4 plots the execution time for the forward and adjoint evaluation of both deltas and vegas, relative to the cost of simply valuing the swaption portfolio. The two curves marked with circles compare the forward and adjoint calculations of all deltas; the curves marked with stars compare the combined calculations of all deltas and vegas.

As expected, the relative cost of the forward method increases linearly with N , whereas the relative cost of the adjoint method is approximately constant. Moreover, adding the vega calculation to the delta calculation substantially increases the time required using the forward method, but

this has virtually no impact on the adjoint method because the deltas and vegas use the same adjoint variables.

It is also interesting to note the actual magnitudes of the costs. For the forward method, the time required for each delta and vega evaluation is approximately 10% and 20%, respectively, of the time required to evaluate the portfolio. This makes the forward method 10–20 times more efficient than using central differences, indicating a clear superiority for forward pathwise evaluation compared with finite differences for applications in which one is interested in the sensitivities of a large number of different financial products. For the adjoint method, the observation is that one can obtain the sensitivity of one financial product (or a portfolio) to any number of input parameters for less than the cost of the original product evaluation.

The reason for the forward and adjoint methods having much lower computational cost than one might expect, relative to the original evaluation, is that in modern microprocessors, division and exponential function evaluation are 10–20 times more costly than multiplication and addition. By reusing quantities such as $L_i(n+1)/L_i(n)$ and $(1 + \delta_i L_i(n))^{-1}$, which have already been evaluated in the original calculation, the forward and adjoint methods can be implemented using only multiplication and addition, making their execution very rapid.

Conclusions

We have shown how an adjoint formulation can be used to accelerate the calculation of Greeks by Monte Carlo simulation using the pathwise method. The adjoint method produces exactly the same value on each simulated path as would be obtained using a forward implementation of the pathwise method, but it rearranges the calculations – working backward along each path – to generate potential computational savings.

The adjoint formulation outperforms a forward implementation in computing the sensitivity of a small number of outputs to a large number of inputs. This applies, for example, in a fixed-income setting, in which the output is the value of a derivatives book and the inputs are points along the forward curve. We have illustrated the use of the adjoint method in the setting of the Libor market model and found it to be fast – very fast. ■

Michael Giles is professor of scientific computing at the Oxford University Computing Laboratory. Paul Glasserman is the Jack R Anderson professor at Columbia Business School. The research was supported in part by NSF grant DMS0410234. Email: giles@comlab.ox.ac.uk, pg20@columbia.edu,

REFERENCES

- Brace A, D Gatarek and M Musiela, 1997**
The market model of interest rate dynamics
Mathematical Finance 7, pages 127–155
- Giles M and N Pierce, 2000**
An introduction to the adjoint approach to design
Flow, Turbulence and Control 65, pages 393–415
- Glasserman P, 2004**
Monte Carlo methods in financial engineering
Springer-Verlag, New York
- Glasserman P and X Zhao, 1999**
Fast Greeks by simulation in forward Libor models
Journal of Computational Finance 3, pages 5–39
- Griewank A, 2000**
Evaluating derivatives: principles and techniques of algorithmic differentiation
Society for Industrial and Applied Mathematics
- Protter P, 1990**
Stochastic integration and differential equations
Springer-Verlag, Berlin