SENSITIVITY ESTIMATION IN FINANCIAL ENGINEERING AND RISK MANAGEMENT

Paul Glasserman

Columbia Business School
403 Uris Hall, Columbia University
New York, NY 10027

This presentation considers the problem of sensitivity estimation through simulation, with a focus on applications in financial engineering and risk management. We review the background of the problem and discuss methods for addressing it. We consider connections between methods and techniques for improving their performance. We also discuss current research directions in the area.

1 BACKGROUND

The problem of calculating price sensitivities is central to both the theory and practice of financial engineering. Through a line of work that begins with the Black-Scholes-Merton framework, we know that, under suitable conditions, a derivative security can be replicated or hedged through trading in other assets; the appropriate trading strategy is characterized by the sensitivity of the price of the security to be hedged with respect to the prices of the assets used to hedge. This principle underlies the operation of the derivatives industry, and it provides the mechanism by which financial institutions manufacture new products and manage their risks.

Indeed, the calculation of price sensitivities is often the primary motivation for developing a pricing model. The prices themselves can sometimes be observed in the market (and will be set by the market regardless of what the models say), but price sensitivities cannot be observed and therfore require modeling and computational methods. With essentially any computational method, the calculation of sensitivities is more difficult than the calculation of prices. This is certainly true of simulation, in which sensitivity estimates can easily require 10 to 100 times more computing time than price estimates to achieve the same level of precision.

In the discrete-event simulation context, the problem of sensivity estimation generated considerable interest beginning in the 1980s, particulary through the work of Ho and Cao (1983), Suri and Zazanis (1988), Glynn (1987), Reiman and Weiss (1989) and Rubinstein (1989). This line of research developed methods that avoid finite-difference approximations to derivatives — that is, they estimate the sensitivity to a model parameter without ever changing the value of that parameter.

These methods fall into two broad categories: methods that differentiate the evolution of a simulated path with all random outcomes held fixed (called *perturbation analysis* or *pathwise* estimates), and methods that differentiate a probability density associated with each outcome (called *likelihood ratio method* or *score function* estimates). The domains of the two types of methods overlap, but each method is applicable to some problems for which the other is not.

These two basic approaches are also applicable in the financial engineering context, where the underlying model is often specified through a stochastic differential equation rather than a discrete-event system. Hedging an option, for example, requires the sensitivity of the option price with respect to the initial price of the underlying asset — this is the option's "delta." A pathwise estimate of delta recursively calculates the sensitivity of the evolution of the underlying asset price with respect to its initial condition, and then calculates the sensitivity of the option payoff to the evolution of the underlying asset. A likelihood ratio method (LRM) estimate differentiates (the logarithm of) the density of the path of the underlying asset with respect to the initial value, and then multiplies this random weight by the option payoff. The pathwise estimate generally requires (at least) continuity of the option payoff, whereas the LRM estimate requires knowledge of a density or transition density associated with the underlying model. Further background on these methods and their application can be found in Chapter 7 of Glasserman (2004). and in the references given there.

2 RECENT DEVELOPMENTS

2.1 Malliavin Estimators

The work of Fournié, Lasry, Lebuchoux, Lions and Touzi (1999) has given rise to a large literature on the use of Malliavin calculus to derive sensitivity estimators for stochastic differential equations. Like LRM estimators, these estimators "differentiate" by multiplying a payoff by a random weight, and thus do not require differentiability of the payoff function. In contrast to LRM estimators, those derived through Malliavin calculus do not require knowledge of a transition density for the underlying model.

The Malliavin estimators are derived directly in a continuous-time model, using powerful tools from Malliavin calculus. When implemented in a simulation, however, they require time-discretization, as does simulation of the underlying model. Time-discretization of the underlying model typically produces a general state space Markov chain, to which the pathwise and LRM estimators are potentially applicable. Thus, the Malliavin route differentiates first and then discretizes, whereas the traditional methods would discretize first and then differentiate. This raises the question of what relation holds between the final estimators.

Chen and Glasserman (2006) investigate this question and the connections among these methods. They show, for example, that a Malliavin delta estimator arises as the continuous-time limit of an average of combined pathwise and LRM estimators. Each of these combined estimators differentiates pathwise up to some point and then differentiates a transition density (which is Gaussian in the discretized model); they differ in the point at which they switch. Chen and Glasserman (2006) establish other relationships for sensitivities to drift and diffusion parameters.

2.2 Adjoint Methods

Pricing models in financial engineering often result in high-dimensional state vectors. This is particularly true in interest rate modeling, which typically describes the evolution of the full term structure of interest rates. This can easily result in state vectors with 20–100 components. For hedging or risk management purposes, one is often interested in sensitivities with respect to many or all of the initial values of this state vector. A straightforward application of the pathwise method to a vector of length M requires propagation of a matrix of $M \times M$ sensitivities: the (i,j) entry of this matrix is the sensitivity of the ith component to the initial value of the jth component. This can add substantial overhead to the calculation of price sensitivities.

Giles and Glasserman (2006) apply an adjoint technique to this problem, drawing on ideas used in computational fluid dynamics, as in Giles and Pierce (2000). The adjoint method produces exactly the same value as the ordinary pathwise estimate on each path, but it rearranges the calculations in a way that can produce these values more quickly. Briefly, the adjoint method is advantageous in calculating sensitivities of a small number of functions (i.e., payoffs) to a large number of parameters; the usual (forward) implementation of the pathwise method is preferable when the number of parameters is small or when sensitivities are required for many different functions.

2.3 LRM From Transforms

As already noted, LRM estimators require knowledge of the transition density of an underlying model. For many models used in financial engineering, a transition density is known only through its characteristic function or Laplace transform. This is the case for many Lévy processes and for the family of affine jump-diffusions (as in Duffie, Pan and Singleton 2000).

In current work, Glasserman and Liu investigate the application of LRM estimators using a characteristic function or Laplace transform. In principle, the most straightforward approach is to apply numerical inversion to the transforms of the density and its derivative; the ratio of these inverse transforms yields the score function, the random weight used by the LRM estimator. We investigate the steps required to carry out, and we analyze the different sources and types of error in the overall procedure. We also investigate the use of approximation techniques as alternatives to numerical transform inversion.

REFERENCES

Chen, N., and P. Glasserman. 2006. Malliavin Greeks without Malliavin calculus. Submitted.

Duffie, D., J. Pan, and K. Singleton. 2000. Transform analysis and option pricing for affine jump-diffusions. *Econometrica* 68:1343–1376.

Fournié, E., J.-M. Lasry, J. Lebuchoux, P.-L. Lions and N. Touzi. 1999. Applications of Malliavin calculus to Monte Carlo methods in finance. *Finance and Stochastics* 3:391–412.

Giles, M., and P. Glasserman. 2006. Smoking adjoints: fast Monte Carlo Greeks. *Risk* 19:88–92.

Giles, M., and N. Pierce. 2000. An introduction to the adjoint approach to design. *Flow, Turbulence and Control* 65:393–415.

Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering.* New York: Springer-Verlag.

Glynn, P.W. 1987. Likelihood ratio gradient estimation: an overview. In *Proceedings of the Winter Simulation Conference.* 366–374. Piscataway, New Jersey: IEEE Press.

Ho, Y.C., and X.-R. Cao. 1983. Optimization and perturbation analysis of queueing networks. *Journal of Optimization Theory and Applications* 40:559–582.

Reiman, M., and A. Weiss. 1989. Sensitivity analysis for simulations via likelihood ratios. *Operations Research* 37:830–844.

Rubinstein, R. 1989. Sensitivity analysis and performance extrapolation for computer simulation models. *Operations Research* 37:72–81.

Suri, R., and M. Zazanis. 1988. Perturbation analysis gives strongly consistent sensitivity estimates for the M/G/1 queue. *Management Science* 34:39–64.