FAST SENSITIVITY COMPUTATIONS FOR MONTE CARLO VALUATION OF PENSION FUNDS

BY

MARK JOSHI AND DAVID PITT

ABSTRACT

Sensitivity analysis, or so-called 'stress-testing', has long been part of the actuarial contribution to pricing, reserving and management of capital levels in both life and non-life assurance. Recent developments in the area of derivatives pricing have seen the application of adjoint methods to the calculation of option price sensitivities including the well-known 'Greeks' or partial derivatives of option prices with respect to model parameters. These methods have been the foundation for efficient and simple calculations of a vast number of sensitivities to model parameters in financial mathematics. This methodology has yet to be applied to actuarial problems in insurance or in pensions. In this paper we consider a model for a defined benefit pension scheme and use adjoint methods to illustrate the sensitivity of fund valuation results to key inputs such as mortality rates, interest rates and levels of salary rate inflation. The method of adjoints is illustrated in the paper and numerical results are presented. Efficient calculation of the sensitivity of key valuation results to model inputs is useful information for practising actuaries as it provides guidance as to the relative ultimate importance of various judgments made in the formation of a liability valuation basis.

KEYWORDS

Actuarial valuation, pensions, adjoints, delta, pathwise method, Monte Carlo.

1. INTRODUCTION

Adjoint methods have recently become popular for computing sensitivities to model parameters when doing derivatives pricing. These methods are an offshoot of the theory of automatic differentiation. This theory shows that if a model has n state variables, then it is possible to compute sensitivities to the initial values of these variables in a fixed finite multiple (which is independent of n) of the time taken to compute the original number. The proof is constructive and relies on decomposing algorithms into simple arithmetic operations which can be trivially, and, in fact, *automatically*, differentiated. We refer the

Astin Bulletin 40(2), 655-667. doi: 10.2143/AST.40.2.2061132

reader to Griewank (2000) for the general theory. In particular, Giles and Glasserman (2006) demonstrated massive speed-ups for the computation of price sensitivities in LIBOR market models for pricing exotic interest rate derivatives.

Adjoint techniques are applicable to sensitivity computation for any algorithm where the outputs vary smoothly with the initial inputs and parameters. Thus far they have not been applied to actuarial applications. Booth et al. (1999) describe the importance of sensitivity analysis: it focuses attention on a range of possible outcomes thus eliminating the possibility of holding too great confidence in the results of a single projection, it informs the user of the level of doubt associated with deterministic projections and it allows the actuary to determine which assumptions are most critical to the outcome of the analysis. Brender (1988) discusses the use of sensitivity analysis in the dynamic solvency testing of life offices. Sensitivity analysis is now part of the requirements for life and non-life assurance offices in their financial condition reports in Australia, the UK and in Canada.

In this paper, we demonstrate the power of adjoints in computing sensitivities by studying the problem of valuing a defined benefit pension fund with interest rate and inflation components. We work in the context of a model described in Neill (1977). In particular, we show that it is possible to compute hundreds of sensitivities and a central estimate of the liability in a couple of seconds, when computing the value alone takes about a second. Actuaries involved in the valuation of pension funds or the pricing of life assurance products will find considerable value in this methodology. The relative magnitudes of sensitivities of calculated outputs, such as premium rates or assurance/pension liabilities to key input assumptions serve as a useful indicator of the relative importance of making accurate assumptions. Drawing on a number of potential actuarial applications of the methodology, sensitivities of valuation results to levels of assumed investment returns, mortality and morbidity rates, pension or assurance scheme withdrawal rates, expense, profit or prudential margins can all be quantified using the method given here.

Our analysis here is focused here on *model* sensitivities, that is sensitivities to the changes of model parameters. We remark that often it is *market* sensitivities that are required. We refer the reader to Joshi-Kwon (2010) for discussion of how one might carry out the translation in a similar context.

This paper is organized as follows. After this introduction, Section 2 outlines the pension model employed in the paper. Section 3 gives the method for finding sensitivities of valuation results to state variables or initial assumptions regarding levels of key economic variables. Section 4 provides the method for assessing valuation result sensitivities to model parameters. Section 5 considers the sensitivity of valuation results to changes in the nature of the benefits provided to members in the pension scheme. Numerical results of our analysis are presented in Section 6. Section 7 compares the method of adjoints with that of finite differencing. We examine the scope of the adjoint method in Section 8. Section 9 concludes the paper.

2. The model

Following Booth et al (1999), we use a discrete-time extended Vasicek process with time-dependent parameters for each of interest rates, i, and inflation, f. We thus have

$$i_{j+1} - i_j = k_{i,j}(i_j - \mu_{i,j}) + \sigma_{i,j} W_j, \tag{2.1}$$

$$f_{j+1} - f_j = k_{f,j}(f_j - \mu_{f,j}) + \sigma_{f,j} Z_j, \tag{2.2}$$

with W_j, Z_j standard normal random variables. We assume that W_j and Z_j are binormal with correlation coefficient ρ and that steps are independent of each other.

The parameters $k_{i,j}$, $\mu_{i,j}$, $\sigma_{i,j}$, $k_{f,j}$, $\mu_{f,j}$ and $\sigma_{f,j}$ can be used to calibrate the model. This results in $6N$ parameters with N the number of time steps from the date of valuation until the death of the pension fund member. We also have the inputs i_0, f_0 and ρ.

The inflation and interest rates are used to evolve the retail price index, RPI, and discretely compounding money market account used for discounting. Thus we have also

$$\mathrm{RPI}_{j+1} = \mathrm{RPI}_j(1 + f_j), \tag{2.3}$$

$$D_{j+1} = \frac{D_j}{1 + i_j}. \tag{2.4}$$

We will take $\mathrm{RPI}_0 = 1$, $D_0 = 1$, since we are interested in relative inflation and discounting. Our model therefore has 4 state-variables. When carrying out valuation, we evolve these variables across N years (with N typically 90) work out the cash-flows and discount these using D_j as they arrive.

3. Adjoint calculations for state-variables

Suppose our product generates a cash-flow each each year which is a function of the prevailing state-variables, as well as possible product parameters, and indirectly model parameters.

First, we discuss how to compute sensitivities to the model variables,

$$X_0 = \left(i_0, f_0, \mathrm{RPI}_0, D_0\right).$$

Whilst in practice, we will not want sensitivities to the last two because these are just benchmark starting values for our price index and our discount factor, we will obtain a more coherent framework by allowing them.

Our discounted cash-flow at time j is some function g_j of X_j. For example, it might be

$$\mathrm{RPI}_j\, D_j\, S_j\, P_j\, C_0,$$

with S_j the survival probability for time j and P_j the payment fraction of salary at time j, and C_0 the initial salary. Our estimate of the value on path s, which we denote by $\mathrm{Value}_s(0)$, is the sum of these discounted cash-flows. Our estimate of the value based on M paths is then

$$\mathrm{Value}(0) = \frac{1}{M}\sum_{s=1}^{M} \mathrm{Value}_s(0). \tag{3.1}$$

Our objective is compute

$$\sum_k \frac{\partial g_k}{\partial X_0}$$

for each path of the simulation. We can write

$$X_j = F_j(X_{j-1}, A_j), \quad \text{for} \quad j > 0, \tag{3.2}$$

with A_j a pair of independent standard normals, and where F_j is the map taking the state-variables from one time to the next. The derivative $\frac{\partial X_j}{\partial X_{j-1}}$ is therefore the Jacobian of the map F_j.

We can therefore write

$$\frac{\partial g_k}{\partial X_0} = \frac{\partial g_k}{\partial X_k}\frac{\partial X_k}{\partial X_{k-1}}\frac{\partial X_{k-1}}{\partial X_{k-2}} \cdots \frac{\partial X_1}{\partial X_0}. \tag{3.3}$$

There are many ways to evaluate this equation which lead to the same final answer but yield vastly differing computational speeds. The first method is simply to compute all the entries of all the Jacobians, multiply the Jacobians together and then multiply by the derivative of g. We then have to carry out k matrix multiplications each of which takes 4^3 multiplications.

A faster method is to multiply from left to right. The associativity of matrix multiplication guarantees that the same answer is obtained but at each stage we are multiplying a vector by a matrix so we have 4^2 multiplications per step instead of 4^3. Thus we set

$$V_k = \frac{\partial g_k}{\partial X_k},$$

and

$$V_{k-j-1} = V_{k-j}\frac{\partial X_{k-j}}{\partial X_{k-j-1}}.$$

For a particular model, by studying the particular structure of F_{k-j} one can do better. For our model, most of the entries of the Jacobian are zero. One can explicitly compute the relationships between the entries of V_{k-j-1} and V_{k-j}. Before doing so, we observe that i_j depends only on i_{j-1}, f_j on f_{j-1}, RPI_j on f_{j-1} and RPI_{j-1}, and D_j on D_{j-1} and i_{j-1}. This means that

$$V_{k-j-1,0} = V_{k-j,0}\frac{\partial i_{k-j}}{\partial i_{k-j-1}} + V_{k-j,2}\frac{\partial \mathrm{RPI}_{k-j}}{\partial i_{k-j-1}}, \tag{3.4}$$

$$V_{k-j-1,1} = V_{k-j,1}\frac{\partial f_{k-j}}{\partial f_{k-j-1}} + V_{k-j,3}\frac{\partial D_{k-j}}{\partial f_{k-j-1}}, \tag{3.5}$$

$$V_{k-j-1,2} = V_{k-j,2}\frac{\partial \mathrm{RPI}_{k-j}}{\partial \mathrm{RPI}_{k-j-1}}, \tag{3.6}$$

$$V_{k-j-1,3} = V_{k-j,3}\frac{\partial D_{k-j}}{\partial D_{k-j-1}}. \tag{3.7}$$

We only have six multiplications per step once the partial derivatives have been computed.

The partial derivatives are all straightforward:

$$\frac{\partial i_m}{\partial i_{m-1}} = (1 + k_{i,m-1}), \tag{3.8}$$

$$\frac{\partial f_m}{\partial f_{m-1}} = (1 + k_{f,m-1}), \tag{3.9}$$

$$\frac{\partial \mathrm{RPI}_m}{\partial \mathrm{RPI}_{m-1}} = (1 + f_{m-1}), \tag{3.10}$$

$$\frac{\partial \mathrm{RPI}_m}{\partial f_{m-1}} = \mathrm{RPI}_{m-1}, \tag{3.11}$$

$$\frac{\partial D_m}{\partial D_{m-1}} = \frac{1}{1 + i_{m-1}}, \tag{3.12}$$

$$\frac{\partial D_m}{\partial i_{m-1}} = -\frac{D_{m-1}}{(1 + i_{m-1})^2}. \tag{3.13}$$

To compute the sensitivities to state variables, on each path we therefore first evolve forward in time computing all the cash-flows generated and storing the values of the partial derivatives of the state variables as we go. We then go back along each path computing V_{N-j} recursively for all j. Since the adjoint calculation for each g_k is the same except for the number of steps, we add on the derivative of g_k when doing step k. Thus we only need to perform the adjoint multiplications once per step for each path.

This yields the derivatives of the discounted cash-flows with respect to the initial values of the state variables for each path. The overall derivative estimate is their average across many paths. In addition to finding the sensitivity to say the initial values of state variables, it is often necessary to analyse the effect of a change in a model parameter, such as an assumed age-specific mortality rate or set of mortality rates, at some future time. Quick calculations of such sensitivities, using the method of adjoints, are given in the following section.

4. Adjoint calculations for model parameters

In section 3, we studied the problem of computing sensitivities to the initial values of the state variables. However, we will also want sensitivities to the model's parameters. In particular, if we allow different parameters for each time step then we have $6N$ parameters. We will see in this section how to compute all these sensitivities with only a small amount of additional effort.

Let θ be a parameter of the evolution of X_{j-1} to X_j. Varying θ will only affect the values of discounted cash-flows which depend on X_l for $l \geq j$. We can write the sensitivity for a given path as

$$W_\theta = \sum_{l \geq j} \frac{\partial g_l}{\partial X_j} \frac{\partial X_j}{\partial \theta}.$$

We can write

$$W_\theta = V_j \frac{\partial X_j}{\partial \theta}.$$

The term V_j has already been computed for the state-variable sensitivity so the additional work is small.

Computation of the terms $\frac{\partial X_j}{\partial \theta}$, are straight-forward. For the coefficients of the extended Vasicek processes, we have

$$\frac{\partial i_j}{\partial k_{i,j-1}} = i_{j-1} - \mu_{i,j-1}, \tag{4.1}$$

$$\frac{\partial i_j}{\partial \mu_{i,j-1}} = -k_{i,j-1}, \tag{4.2}$$

$$\frac{\partial i_j}{\partial \sigma_{i,j-1}} = W_{j-1}, \tag{4.3}$$

$$\frac{\partial f_j}{\partial k_{f,j-1}} = f_{j-1} - \mu_{f,j-1}, \tag{4.4}$$

$$\frac{\partial f_j}{\partial \mu_{f,j-1}} = -k_{f,j-1}, \tag{4.5}$$

$$\frac{\partial f_j}{\partial \sigma_{f,j-1}} = Z_{j-1}. \tag{4.6}$$

Note that RPI_j and D_j do not have direct parameter dependence, since they are evolved as functions of their previous value and i_j and f_j. They do, of course, have indirect dependence via the effects of parameters on i_j and f_j.

The dependence on ρ is a little more interesting. Our random number generator will yield two independent random normals per step, $A_{1,j}$ and $A_{2,j}$. We typically set

$$W_j = A_{1,j}, \tag{4.7}$$

$$Z_j = \rho A_{1,j} + \sqrt{1-\rho^2}\, A_{2,j}. \tag{4.8}$$

We then have that the evolution of i_j does not depend on ρ, but that the evolution of f_j does. In particular,

$$\frac{\partial Z_j}{\partial \rho} = A_{1,j} - \frac{\rho}{\sqrt{1-\rho^2}} A_{2,j}.$$

And so,

$$\frac{\partial f_j}{\partial \rho} = \sigma_{f,j-1}\left(A_{1,j} - \frac{\rho}{\sqrt{1-\rho^2}} A_{2,j}\right). \tag{4.9}$$

Since we are using the same ρ parameter for all steps, we obtain that the overall ρ sensitivity is

$$\sum_j V_j \frac{\partial X_j}{\partial \rho}.$$

So far our analysis has produced sensitivities of valuation results to changes in the values of state variables and to changes in the values of model parameters. We can also investigate how our valuation result changes when more basic product design changes are made. The calculation of these sensitivities, using the method of adjoints, is discussed in the next section.

5. Pay-off sensitivities

As well considering sensitivities to changes in our model state variables and parameters, we may also wish to know how changing the specification of our product changes its value. In this section, we address the problem of computing sensitivities to the pay-off specification. If we are using deterministic survival probabilities, it is also more effective from a computational perspective to consider them as part of the product specification. Thus our product pays a sequence of discounted cash-flows (possibly positive or negative),

$$g_l(X_l, \phi)$$

and we want to know the derivatives of the expected pay-off with respect to ϕ. In fact, provided the functions g_l are Lipschitz continuous we can simply differentiate the pay-offs and compute

$$\sum_{l=1}^{N} \frac{\partial g_l}{\partial \phi_r}$$

for each path and average.

For example, if

$$g_l = \mathrm{RPI}_l \, D_l \, S_l \, P_l \, C_0,$$

then we would have

$$\frac{\partial g_l}{\partial C_0} = \mathrm{RPI}_l \, D_l \, S_l \, P_l, \tag{5.1}$$

$$\frac{\partial g_l}{\partial S_k} = \delta_{l,k} \, \mathrm{RPI}_l \, D_l \, P_l \, C_0, \tag{5.2}$$

$$\frac{\partial g_l}{\partial P_k} = \delta_{l,k} \, \mathrm{RPI}_l \, D_l \, S_l \, C_0, \tag{5.3}$$

where $\delta_{l,k}$ equals 1 if $l = k$, and zero otherwise.

6. Example

Consider a defined benefit superannuation fund. Our model for the value of this fund follows those described in Neill (1977). The member joins at age 20 and remains a contributing member to the fund until age 60 or earlier death. Contributions to the fund are a percentage of annual salary taken for this example to be 5%. Beyond age 60, the member withdraws money from the fund equal to some proportion of the member's final average salary each year contingent on survival. This salary withdrawal rate is again taken to be 5% per annum for this example. Salary increases in line with inflation where inflation is modeled using (2.2). We have chosen $f_0 = 0.05$, $\mu_{f,j} = 0.05$, $k_{f,j} = -0.3$ and $\sigma_{f,j} = 0.01$. Survival probabilities for the member are set equal to those in the 2000-2002 Australian life tables. The cash inflows and outflows from the fund are discounted back to the fund opening date using an annual effective rate of interest modeled using (2.1). We have chosen $i_0 = 0.05$, $\mu_{i,j} = 0.05$, $k_{i,j} = -0.3$ and $\sigma_{i,j} = 0.01$. This present value is a measure of the financial viability of the fund as it indicates the ability of the fund to meet its future financial obligations. We study the sensitivity of this present value to changes in: the current level of interest rates, the current level of inflation, the annual salary contribution percentage made during the working life of the member, the percentage of final average salary withdrawn from the fund during retirement and a decrease in the annual force of mortality across all ages. These sensitivities are either an immediate output of the methodology presented in Sections 3, 4 and 5 or are obtained as simple linear combinations of sensitivity outputs obtained in those sections.

The results of this sensitivity analysis are given in Table 1. Immediate from these results is that assumptions relating to salary, both rate of contribution and rate withdrawn during retirement have the most significant effect on the valuation result. The valuation result here is positive when we project that the retirement benefit commitments of the fund will be able to be met by the accumulated value of salary contributions. We note also that the impact of higher investment returns leads to an increased ability of the fund to meet its commitments, however the sensitivity of the valuation result in this instance is materially lower than in the case of salary rate contributions. Positive inflation, which is applied to the salary levels in our model, affects the ability of the fund to meet its commitments adversely, as shown by the negative sensitivity result in Table 1. Finally, the impact of a reduction in the mortality rate levels across all ages reduces the ability of the fund to meet its commitments due to the resulting longer period of retirement benefit payments. Mortality rates improvement of 1% has the lowest impact on our valuation result. It would be worth noting that mortality rates have been shown, using empirical studies,

TABLE 1

PARAMETER SENSITIVITIES FOR THE PENSION FUND VALUE

Change to Actuarial Basis	**Sensitivity**	**FD sensitivity**
Interest Rates: +1%	–0.0223	–0.0223
Inflation Rate: +1%	0.0223	0.0223
Salary contribution: +1%	0.3927	0.3927
Force of mortality: –1%	–0.0042	–0.0042
Salary withdrawal: +1%	–0.2099	–0.2099

for example Lee and Carter (1992), to have different rates of improvement at different ages. The sensitivity of a valuation result to differing mortality rate improvement levels at different ages can be calculated using the adjoint methodology given here.

In Table 1, we also present a comparison against numbers computed using finite differencing labelled “FD”. As expected, these numbers are identical showing that the use of the adjoint method has caused no change in accuracy. We will discuss this comparison further in Section 7.

Results providing information on the relative importance of assumptions and other model parameters can be found very quickly. The sensitivity of valuation results to multiple changes in the model parameters or actuarial basis can be determined as linear combinations of the sensitivities reported above.

7. Adjoints versus finite differencing

The obvious naive method to compute sensitivities is to bump a parameter a small amount, recompute the price and compute the finite difference sum. Thus if our parameter is θ, we compute

$$\frac{\mathrm{Value}(0,\theta+h)-\mathrm{Value}(0,\theta)}{h}.$$

If the estimate of the value on path s is $\mathrm{Value}_s(0,\theta)$ then we can write our finite difference estimate of the sensitivity using M paths as

$$\frac{1}{Mh}\left(\sum_{s=1}^{M}\mathrm{Value}_s(0,\theta+h)-\sum_{s=1}^{M}\mathrm{Value}_s(0,\theta)\right),$$

which equals

$$\sum_{s=1}^{M}\frac{1}{Mh}\big(\mathrm{Value}_s(0,\theta+h)-\mathrm{Value}_s(0,\theta)\big).$$

Provided the $\mathrm{Value}_s(0,\theta)$, is twice continuously differentiable as a function of θ, then this is equal to O

$$\frac{1}{M}\sum_{s=1}^{M}\frac{\partial\mathrm{Value}_s}{\partial\theta}(0,\theta)+O(h).$$

For h small this means that the finite difference estimate will agree with the path-wise adjoint method on every path except for a very small discretization bias. This has two immediate consequences. First, when testing the implementation the fact that the two methods agree on every path means that errors are

TABLE 2

TIME TAKEN TO COMPUTE PRICE AND ALL SENSITIVITIES AS A FUNCTION OF THE NUMBER OF YEARS

N	Time taken
10	0.252
20	0.5
30	0.688
40	0.936
50	1.124
60	1.44
70	1.624
80	1.816
90	2.064

easy to discover. Second, the standard error of the Monte Carlo simulation for estimating the path-wise sensitivity will have almost the same standard error as that for the finite differencing. We therefore do not lose (or gain) anything in terms of convergence rate as a function of the number of paths by using the adjoint method.

In fact, the path-wise method can be applied even when the pay-off is not twice-differentiable. The crucial fact required is that the evaluation of the pay-off is Lipschitz continuous in the parameter to be differentiated. We refer the reader to Glasserman (2004), Section 7.2.2 for further details.

When using finite differencing, however, we must run an extra simulation for every sensitivity, and two simulations if we choose to use centred differencing to minimize discretization bias. Thus to compute 720 sensitivities it will take 720 times as long as computing the price.

With the adjoint method, the time increase is very small by comparison. The reason being that the recursive computations used as detailed in Section 3 are very simple, requiring only an extra 18 multiplications for each step of very path to compute the state variable sensitivities. For the parameter sensitivities, each one only requires a small number of additional floating point operations per path.

For example, a pension was priced using 2^{16} paths with $N = 90$. The simulation was implemented using single-threaded C++ on a Quad-Core Xeon. The time taken to obtain the price was 1.32 seconds. To obtain the price and 724 sensitivity numbers took 2.06 seconds. Thus computing all possible sensitivities took less time than computing one sensitivity using finite differencing. We present the time taken to compute all sensitivities as a function of the number of years with 2^{16} paths in Table 2. We see that the time taken is approximately linear in the number of time periods.

We do not present comparisons of the standard errors, since as noted above they would be almost identical.

8. Scope of the method

We have considered a case where the state variables evolve via normal increments and the cash-flows are continuous functions of these variables. We now discuss how generally the method applies. The crucial assumption is that the evolution of the state variables is given by a differentiable function of the previous variables and parameters of interest. Thus we must have a sequence of maps

$$F_j : X_{j-1} \mapsto X_j$$

such that we can write

$$X_j = F_j(X_{j-1}, \theta_j, Z_j),$$

where θ_j are the parameters of interest and Z_j are random variables drawn from some fixed distribution, with F_j a differentiable function of X_j and θ_j.

The results in Griewank (2000) then guarantee that is possible to carry out the adjoint computation in a fixed finite multiple of the original time to compute the value of the fund. In practice, one would decompose the maps F_j into simpler maps until the adjoint calculation was simple enough to make this straightforward.

For example, it would be a simple extension to add an extra state-variable representing the value of an investment fund following a process such as

$$\log S_{j+1} = \log S_j + a_j i_j + \beta_j f_j + \gamma_j + \sigma_j Q_j$$

with Q_j a normal random variable (or, indeed, any random variable from a fixed distribution).

In practice, F_j may represent an approximation to the true evolution; we are differentiating a discretization of the model, not the mathematical model itself. This rules out simulation techniques where a small change in parameters or variables can causes a jump in value. For example, acceptance-rejection methods are not compatible with this approach.

Our secondary requirement is that the cash-flows be Lipschitz continuous: a cash-flow which jumps in value will break the method. The method can be extended to such cases; however, the theory becomes considerably more involved. We refer the reader to Chan-Joshi (2009) for a discussion of such an extension in a different context.

9. Conclusion

This paper has applied the method of adjoints, recently used for very efficient calculations of sensitivities of derivatives prices, to assessing sensitivities of valuation results for pension funds. The method is extremely fast computationally and enables actuaries to focus more of their time on the actuarial judgments associated with selecting a valuation basis in life and non-life assurance as well as pensions, than on computations. We have shown that the method quickly produces a vast number of sensitivities which can be combined in many ways to assess the effect on valuation results of a wide range of possible scenarios.

References

BOOTH, P., CHADBURN, R., COOPER, D., HABERMAN, S. and JAMES, D. (1999) *Modern Actuarial Theory and Practice.* Chapman and Hall/CRC, London.

BRENDER, A. (1988) *Testing the solvency of life insurers*. Proceedings of the 2nd International Conference on Insurance Solvency.

CHAN, J.-H. and JOSHI, M. (2009) *Fast Monte-Carlo Greeks for Financial Products With Discontinuous Pay-Offs*. http://ssrn.com/abstract=15003

GLASSERMAN, P. and GILES, M. (2006) *Smoking adjoints: fast Monte Carlo Greeks*. Risk, January 2006, 92-96.

GLASSERMAN, P. (2004) *Monte Carlo Methods in Financing Engineering.* Springer-Verlag, Berlin-Heidelberg-New York.

GRIEWANK, A. (2000) *Evaluating derivatives: principles and techniques of algorithmic differentiation.* Society for Industrial and Applied Mathematics.

JOSHI, M. and KWON, O.K. (2010) *Monte Carlo Market Greeks in the Displaced Diffusion LIBOR Market Model*. http://ssrn.com/abstract=1535058

LEE, R.D. and CARTER, L. (1992) Modeling and Forecasting the Time Series of US Mortality. *Journal of the American Statistical Association* **87**, 659-671.

NEILL, A. (1977) *Life Contingencies*, Butterworth-Heinemann Ltd.

MARK JOSHI
University of Melbourne
E-Mail: mark@markjoshi.com

DAVID PITT
Department of Actuarial Studies
Macquarie University
NSW 2109
Australia
E-Mail: david.pitt@mq.edu.au