CENTRE FOR EMEA BANKING, FINANCE & ECONOMICS

Local Volatility Calibration with the Markov Chain Approximation

Chia Chun Lo and Konstantinos Skindilias

Working Paper Series

No 29/11

Local Volatility Calibration with the Markov Chain Approximation

Lo Chia Chun and Konstantinos Skindilias

This Version: May 2011

Department of Finance, University of Macau, Macau
London Metropolitan Business School, London Metropolitan University, London, UK

Abstract

We propose a method for calibrating the local volatility surface that relaxes the computational complexity associated with many models and prices options consistently with the volatility skew. We achieve this under a continuous time semi-closed form solution based on the Markov chain approximation of Kushner (1990). The continuous time Markov chain setting can accommodate a diffusion or a jump-diffusion. We show that under both cases without the use of any regularisation we can accurately price out-of-sample options.

Key Words: Local Volatility, Inverse Problem, Markov chain approximation, Spline estimation

1 Introduction

With a simple extension to the Black & Scholes (1973) world, the local volatility model of Dupire (1994) can achieve consistency with the implied volatility while still retain the properties of a complete market. Further, as the well documented addition of a jump to a diffusion creates steep realistic short-term skews, Andersen & Andreasen (2000) extend the local volatility method to incorporate jumps. Although this methodology loses its market completeness, Andersen & Andreasen (2000) show that coupling local volatility and a jump-diffusion allows modeling of the steep short term volatility skew and accurate price fitting.

Following the seminal works of Derman & Kani (1994), Dupire (1994) and Rubinstein (1994)[1] numerous authors have proposed methodologies that effectively encounter the natural ill-posedness of fitting the smile due to the under-specified case. Lagnado & Osher (1997*b*) presented a regularised minimisation method, using finite differences, to fit option prices. Their method forces well-posedness

[1]The authors have independently shown how a class of index option that exhibits an implied volatility smile and/or term structure can be consistently priced using a complete model with a deterministic volatility function of spot and time

of the inverse problem by restricting volatility to the smoothest function possible. It is an effective way but computational expensive. A drawback is the discrete representation of the instantaneous volatility function by a small array of nodes in space and time imposed to reduce computational time. Coleman, Li & Verma (1999) and Jackson, Suli & Howison (1999) use a space-time spline to represent the local volatility function. The lack of sufficient market option price data and time consuming computation is therefore overcome by introducing smoothness in the local volatility function. Option prices are computed numerically using a finite difference approximation. An alternative to these approaches that relies on probabilistic convergence proofs is given in Crepey (2003). Here the local volatility surface is obtained using Tikhonov regularisation under a trinomial tree discretisation. The jump-diffusion case introduced in Andersen & Andreasen (2000) and further examined in He, Kennedy, Coleman, Forsyth, Li & Vetzal (2006) use a forward partial integro differential equation (PIDE) to describe the evolution of European call options.

As for any minimisation method an algorithm is required to price the model options, in this paper we make use of the well-established idea of local consistency given in Kushner & Dupuis (2001). For this purpose we set up a Markov chain in continuous time with a finite number of states. As opposed to the numerical approximations mentioned above, where one needs to discretise time and space, a Markov chain can be viewed as a lattice set in continuous time. The advantage of a continuous time Markov chain lies in the fact that a purely probabilistic approach is undertaken where one only needs to consider discretisation of the state space. For the case of a jump-diffusion Kushner & DiMasi (1978) have shown that the local consistency idea holds. In Chia & Skindilias (2011) it is shown that when the underlying follows a jump-diffusion the Markov chain approximation prices accurately European options. As a result, rather than having to treat the integral form explicitly while the differential terms implicitly (Zhang 1997), or use the Crank-Nicolson method (for the differentials) with adaptive time steps and FFT for evaluating the integral term (Andersen & Andreasen 2000), required by PIDE's, the Markov chain approximation is much simpler to handle. (Given that in a jump-diffusion more states need to be reachable, the Markov chain resembles a multinomial tree.)

The use of the Markov chain approximation in finance is not a novel approach. Application to optimal control problems is given in Kushner (1997). Monoyios (2004) sets up a Markov chain in discrete time to price options with transaction costs, while Chourdakis (2004) examines stochastic volatility models of non-affine nature.

(In this paper we relax the computational complexity of modelling the local volatility surface and show that through a Markov chain approximation we can acccurately fit option prices to the volatility skew. Without the use of regularisation we are able to calibrate a unique volatility surface). In the next section we describe the inverse problem.

2 Local Volatility Calibration

To price any European option value from the smile, one only needs to interpolate the market-observed implied volatilities to the strike and maturity of that option. To compute an out-of-sample price or to price exotics, one then needs to compute the local volatility and use with a numerical method to solve the Black-Scholes partial differential equation (PDE). From the interpolated volatility surface, using the Dupire (1994) equation one can then compute the local volatility surface at each point in the underlying S and time t. Since local volatility varies twice as fast with spot as implied volatility varies with strike, the resulting local volatility surface (and hence option value) will heavily rely on the interpolation scheme used.

Since for the local volatility model the instantaneous volatility is deterministic in spot and time (i.e. $\sigma^2(S_T, T, \cdot) \stackrel{def}{=} \sigma^2(S_T, T)$), instantaneous volatility evolves along the static local volatility function. Therefore to avoid dependence on the interpolation of the implied volatility, the methods of Lagnado & Osher (1997*a*), Coleman et al. (1999), Jackson et al. (1999), Andersen & Andreasen (2000), He et al. (2006) back out the deterministic instantaneous volatility (that defines the dynamics of the PDE) directly from market-observed option prices. This type of inverse method minimises the variational least squares problem:

$$\min_{\sigma(S_t,t)\in\mathcal{H}} \sum_{\forall j} (V_j(\sigma(S_t, t)) - \bar{V}_j)^2. \tag{1}$$

Furthermore, to handle the computational complexity of Lagnado & Osher (1997*a*) where the volatility is evaluated at every point of $\sigma(S_t, t)$, Coleman et al. (1999) and Jackson et al. (1999) choose to represent the local volatility function by a space-time spline. Using splines to represent the local volatility essentially treats the under-specified case of the inverse problem due to finite data.

In practice a spline based representation of the local volatility function will have p number of knots in $\{(S_i, t_i)\}_{i=1}^{p}$ over a finite grid. An interpolating cubic spline $\varsigma(S, t)$ will completely determine the local volatility over the grid at every point of the grid. Given $\sigma'(S_i, t_i)$, the volatilities at the spline knots, the inverse problem of (1) can be formalised as:

$$\min_{\sigma'\in\mathbb{R}^p} \sum_{\forall j} (V_j(\varsigma(S_t, t)) - \bar{V}_j)^2, \tag{2}$$

where $V_j(\varsigma(S_t, t; \sigma')\forall j$ is the option price computed using the instantaneous volatility function.

An option price under the risk neutral probability measure $\mathbb{Q}$ is the expectation (for some payoff function $H(S(T))$) of

$$V_t = e^{-r(T-t)} E_t^{\mathbb{Q}}[H(S(T))]. \tag{3}$$

Rather than computing option prices using a numerical approximation to the PDE or use any method that requires us to discretise time in this paper we

use a continuous time Markov chain to approximate the underlying process. In particular we focus on the local consistency condition of Kushner (1990) that states to maintain weak convergence to the underlying pilot process one only needs to match the first two instantaneous moments of the increments of the approximating Markov chain with those of the continuous-time limiting stochastic differential equation (SDE).

3 The Markov chain approximation with deterministic volatility

Our rather heuristic approach to approximate the local volatility function sets up a continuous time Markov chain with a finite state space $G = [S_{min}, \ldots, S_{max}]$ over the underlying price S. The Markov chain will have a rate generator matrix $Q(t)$ that fully defines the transition probability matrix through the Chapman-Kolmogorov relationship: $P(t) = \exp\{tQ\}$. For many pricing applications the transition kernel is piecewise constant between dates $t_1 = 0, \ldots, t_N = T_N$[2], so that $Q_n = Q(t)$ for $t \in [t_{n-1}, t_n)$. The transition probability matrix is then given by the product of the matrix exponentials

$$P_n = \prod_{i=1}^{N} \exp\{-(t_i - t_{i-1}) \times Q_i\}. \tag{4}$$

As a result of the local volatility function, each generator matrix $Q(t_i)$ (and transition probability matrix) with have state-time dependent elements on the local volatilities at $\sigma(S_{\forall S \in G}, t_i)$. The simplest spline functional will have spline knots positioned on a rectangular grid at positions $\{(S_v, t_m)\}_{v=1,m=1}^{k_s,k_t}$, where the total number of spline knots is $p = k_s \times k_t$. The construction of each transition matrix $P(t)$ will depend on the interpolated volatilities across all states S for a given date $\varsigma(S, t)$. So far we have assumed that the process underlying S has a deterministic volatility function in spot and time, and we approximate this process by the continuous time Markov chain construction described above. As the European option price in (3) is an expectation, computation using Markov chains is straight forward. A European option price with maturity $T_n < T_N$ [3] is simply given by

$$V_t = e^{-r(T-t)} E^{\mathbb{Q}}[H(S(T))] = e^{-r(T-t)} \sum_j P_{i,j}(t) \cdot H(S_j(T)), \tag{5}$$

where $P(t)$ is given by (A)[4].

In a 1-factor pure diffusion process the local volatility of Dupire (1994) is given by

$$dx(t) = \mu^{\mathbb{Q}}(x, t)dt + \sigma(x, t)dW^{\mathbb{Q}}(t)$$

[2] where T_N is the maximum maturity in the given trading day

[3] we will change this to something better. I want to write it in a way so that the reader understands that when he wants to evaluate an option with maturity $T_n < T_N$ he just needs to evaluate (4) (up to time T_n) and the simply use (5)

[4] Maybe it will be smarter to write the price as a function of Arrow Prices since arrow price can represent the RND (or better the SPD) which one can directly use to price an option in future time

The jump-diffusion with a deterministic volatility function of Andersen & Andreasen (2000) takes the form

$$dx(t) = \mu^{\mathbb{Q}}_{JD}(x,t)dt + \sigma(x,t)dW^{\mathbb{Q}}(t) + y(x,t)dP^{\mathbb{Q}}(t),$$

where the discontinuous jump process y can be defined by a Poisson random measure $ydP(t) = \int_{\mathcal{Q}} y(q)\mathcal{P}(dt,dq)$ with $E[\mathcal{P}(dt,dq;x(t),t)|x(t) = x] = \phi_Q(q;x,t)dq\lambda(t;x,t)dt$. Our notation above implies that the $\phi_Q(q;x,t)dq$ probability density can have a state-dependent Poisson amplitude mark with a state-time dependent jump intensity λ, a flexibility of the Markov chain approximation that is hard to incorporate in other models.

4 Practical Implementation

4.1 Data Description

The Markov chain approximation with deterministic volatility is based on European style exercise options on the FTSE-100 index (ESX). For the study end-of-day prices are used from the NYSE-EURONEXT database. In particular we calibrate the local volatility function each Friday and use the estimated volatility and create in-sample prices as well as 3/4/5 day out-of-sample prices (that is Monday/Tuesday/Wednesday) for options traded on the specific dates. This repetitive procedure takes place for the periods between MM/YY until MM/YY. As a result we calibrate the local volatility XXX times for the available XXX Fridays in the dataset. For discounting we use the London Interbank Offer Rate (LIBOR) closest to the relevant option maturity provided by the Bank of England.

Trading volume for FTSE-100 ESX options has a substantial difference for SPX options in the US market. As a result options used in the calibration rely heavily on the trading volume. We follow some filtering criteria of a study on FTSE-100 options from Lin, Strong & Xu (2001). We also follow basic no-arbitrage conditions as also noted in Bakshi, Cao & Chen (1997)[5]. In particular we remove options with absolute moneyness ($M = \frac{K}{S} - 1$) greater than 12%[6], options with less than 6 days to maturity and a price less than £3. FTSE-100 (ESX) options are quoted based on the FTSE-100 futures price, so we extract the dividend from the relationship $F = Se^{(r-q)*(T-t)}$.

4.2 Calibration routines

As also in Coleman et al. (1999),Jackson et al. (1999) and He et al. (2006) the problem here is an inverse minimisation problem with respect to the local volatilities $\sigma'(S,t)$ at the spline knots and any other possible parameters ϑ that the underlying process might depend on. Given the above we define the Markov chain option price as a function of the local volatility surface and other possible parameters as $V(S0,0;\varsigma(S,t;\sigma'),\vartheta,K,T) = V(\varsigma(S,t;\sigma'),\vartheta)$. The local

[5] and the no-arbitrage conditions for a finite options system of Mercurio XXXX

[6] or $M = \frac{K}{S}\sqrt{(T)} - 1$ higher than 1.2

volatility function that is fully determined by its values at the spline knots, and other possible parameters that the underlying process may or may not depend on, can be found by solving the inverse minimisation problem

$$\min_{\sigma',\vartheta} \sum_j \left(V_j(\varsigma(x,t;\sigma'),\vartheta) - \bar{V} \right)^2$$

For the 1-factor diffusion model with deterministic local volatility of Dupire (1994) we use the transformed log-asset price process,

$$dx(t) = (r - \frac{1}{2}\sigma^2(x,t))dt + \sigma(x,t)dW^{\mathbb{Q}}(t),$$

where $x = \ln S$ and r is the interest rate. The inverse minimisation is a inverse spline local volatility problem with respect to the volatility at the knots. Given p spline knots the inverse problem is

$$\min_{\sigma' \in \mathbf{R}^p} \sum_j \left(V_j(\varsigma(x,t;\sigma')) - \bar{V} \right)^2$$

Using a spline functional greatly reduces computational complexity and addresses the under-specified case of calibrating the entire local volatility from a finite set (as in Lagnado & Osher (1997*a*)). As it has been noted mainly in Wahba (1990) the position and number of knots is a complicated issue that affects the functional form of the surface. We do note that in the method proposed in this paper usually 3 knots in space and 3 knots in time placed uniformly over the grid serves our purposes for estimating a unique local volatility. The spline functional has flexibility and one may choose to place time knots on the corresponding maturities. In our case we practically use 5 in space, but only the 3 knots are directly a part of the minimisation process as we set the values at the first and last knots in space equal to their neighboring knots, that is $\sigma'(x_1) = \sigma'(x_1 \times 1.15)$ and $\sigma'(x_{max}) = \sigma'(x_{max} \times .85)$. Generally, minimisation takes place with the least specifications of only setting the spline knots as above and by selecting the grid space. In this case a grid of only 41 points is used that spanning $x = \{-.2 : .01 : .2\}$.

For the case of a jump-diffusion we consider that of Merton (1976). The transformed log-asset price with deterministic volatility of Merton's jump-diffusion takes the form

$$dx(t) = (r - \frac{1}{2}\sigma^2(x,t) - \lambda^{\mathbb{Q}}\kappa^{\mathbb{Q}})dt + \sigma(x,t)dW^{\mathbb{Q}}(t) + \ln(y(t,Q))dP^{\mathbb{Q}}(t;Q)$$

In addition to (?), $dP^{\mathbb{Q}}(t;Q)$ is a Poisson process with $\mathbb{E}[dP(t;Q)] = \lambda(t)dt$ and Q_i denotes the i-th mark of the jump amplitude density. $\lambda^{\mathbb{Q}}$ the jump intensity and $\kappa^{\mathbb{Q}} = E^{\mathbb{Q}}(y-1)$. The jump $\ln(y)$ is a normally distributed random variable with mean $\mu^{\mathbb{Q}}$ and standard deviation $\delta^{\mathbb{Q}}$, such that $\mathbb{E}(y) = e^{\mu^{\mathbb{Q}} + \frac{1}{2}(\delta^{\mathbb{Q}})^2}$. For the Merton jump-diffusion this minimisation problem is an inverse problem that takes the form

$$\min_{\sigma' \in \mathbf{R}^p, \vartheta \in \mathbf{R}^3} \sum_j \left(V_j(\varsigma(x,t;\sigma'),\vartheta) - \bar{V} \right)^2$$

subject to lower and upper bounds on the parameters $b_{lower} \geq \{\sigma', \vartheta\} \geq b_{upper}$. The parameter set $\vartheta = \{\lambda^{\mathbb{Q}}, \mu^{\mathbb{Q}}, \delta^{\mathbb{Q}}\}$ define the parameters of the jump part intensity and amplitude under the risk neutral measure.

The ill-posedness of calibrating the local volatility surface and jump parameters in the Merton jump-diffusion has been extensively discussed in He et al. (2006). The authors note that a unique solution for ϑ is computationally expensive and that calibration of the jump parameters is even ill-posed in the constant volatility Merton jump-diffusion. They further note that the parameters that are mostly responsible for the ill-posedness are the jump mean $\mu^{\mathbb{Q}}$ and standard deviation $\delta^{\mathbb{Q}}$. Nevertheless the authors note that ill-posedness of the jump parameters still produces a unique local volatility function. They also note that to tackle the effect of one of the most responsible parameter, $\delta^{\mathbb{Q}}$, one may also impose some restrictions on the upper/lower bounds of the jump amplitudes standard deviation based on historical information since $\delta^{\mathbb{P}} = \delta^{\mathbb{Q}}$[7].

When constructing a Markov chain for the purpose above it is important to make sure that the grid can accommodate a possible large jump amplitude. (We make sure that the grid spans the appropriate length by first calibrating the Merton (1976) close form equation[8] and examining the estimated jump mean and variance[9].) In this respect for the jump-diffusion applications the Markov chain grid spans the interval $x = \{-.5 : 0.01 : .5\}$. In estimating the local volatility function we use the same specifications for the spline knots as for the pure-diffusion case.

Our benchmark models for in-sample and out-of-sample pricing are the well documented models of Dumas, Fleming & Whaley (1998) and Heston (1993). According to Christoferson (XXX), the practitioners Black-Scholes (PBS) which uses the parametric implied volatility function of Dumas et al. (1998) enjoys the reputation of the best performing model for out-of-sample pricing (for up to at least 5 days). Following the route to testing our pricing errors described in Christoferson we show in our findings below that the MCA with a deterministic volatility and a jump component has great potential to outperforming the PBS for out-of-sample pricing. Both proposed models outperform the benchmarks for in-sample pricing. We measure performance by looking at the root mean square error (RMSE) given by

$$RMSE = \sqrt{(}\frac{1}{M \times N}\sum_{j}(C_{model} - C_{obs}$$

4.3 Calibration results

Under the minimum requirements one needs to set up the approximating chains we first examine in-sample pricing performance and compare with the PBS and Heston's stochastic volatility model. Table X reports RMSE's of all models. The two applications associated with the MCA are denoted as MCA_{DIF} for the pure-diffusion and MCA_{JDIF} for the jump-diffusion. In most calibrations the

[7] A technique to estimate $\delta^{\mathbb{P}}$ is given in Hanson & Westman (2004)

[8] jump parameters ϑ and a constant volatility $\bar{\sigma}$

[9] Maybe remove this

MCA outperform for in-sample the PBS and Heston model. Further, $MCAJ_JD$ outperforms MCA_D in all cases. Heston stochastic volatility model seems to suffer from largest errors.

For the pure-diffusion case the MCA calibration is straight forward giving an acceptable RMSE. Fitting the volatility with splines works well for in-sample pricing but this does not mean that spline placement will not affect the results in any manner. The choice for the position of the spline knots comes after numerous experimentations but it is this unique selection that has been used in all experiments.

While the MCA prices have a very satisfactory in-sample performance, when it comes to out-of-sample pricing PBS is the model with the lowest RMSE. The second best performing model is the MCA_JD with very small differences from the PBS. Heston's model is the model with the worst performance for our samples and application. (This is where to do some work on. I believe that JD can outperform).

Callibration of the jump frequency and amplitude parameters for the MCA_JD model is satisfactory in the sense that all of our estimates have some form of consistency with Merton's constant volatility closed form price. This argument may be based on the fact that jump-parameter estimation does not affect local volatility estimation(He et al. 2006) hence constant or local volatility calibration should not impact jump parameter estimation and vice versa.

	in-sample	4-day	5-day	6-day
MCA_{DIF}	1.1013	5.3602	5.2087	4.5606
PBS	1.6993	3.8500	4.4918	3.0847
$HESTON$	2.0753	11.5808	12.5136	11.0501
MCA_{JDIF}	0.3850	3.8718	4.9306	4.2425

Table 1: In-sample and out-of-sample RMSE's for the four competing models.

It is logical to assume that if deep in-the-money and/or deep out-of-the-money will make calibration easier as more information on the tails of the distribution will be provided. Therefore we run the same calibrations but this time we allow more deep-OTM and deep-ITM options. TBC

Appendices

A MCA

For the Markov chain $\{x^*\}$ to be locally consistent (for a detailed discussion see Kushner & Dupuis 2001) with the underlying process $\{x\}$ one only needs to ensure that over a small infinitesimal interval δ and as the grid spacing $h \to 0$:

$$\begin{aligned} E_t^*\{x^*(t+\delta) - x^*(t)\} &= E_t\{x(t+\delta) - x(t)\} + o(h) \\ E_t^*\{(x^*(t+\delta) - x^*(t))^2\} &= E_t\{(x(t+\delta) - x(t))^2\} + o(h) \end{aligned}$$

The elements of the rate generator matrix $Q(t)$ that will satisfy weak convergence to the underlying process are given by Piccioni (1987)[10]:

$$\begin{aligned} q_{i-1i} &= \frac{1}{2h^2}\sigma^2(x_i,t) + \frac{1}{h}\mu^-(x_i,t) \\ q_{ii} &= -\frac{1}{h^2}\sigma^2(x_i,t) - \frac{1}{h}|\mu(x_i,t)| \\ q_{i+1i} &= \frac{1}{2h^2}\sigma^2(x_i,t) + \frac{1}{h}\mu^+(x_i,t), \end{aligned} \tag{6}$$

where $q_{ij} = 0$ for all other $j \neq i, i-1, i+1$, while $|\mu| =$.

B MCAJUMP

Kushner & DiMasi (1978)[11] show that the Markov chain approximation will be locally consistent (in the sense of (**??**)) with the jump-diffusion as $h \to 0$. Due to the separability of a jump-diffusion process the Q_{JD} rate generator matrix can be represented as $Q = Q_{cont} + Q_{jump}$, where Q_{cont} is given by (A). Chia & Skindilias (2011) show a good choice for Q_{jump} as a discretisation of the density $\mathcal{Q}$ to be

$$\begin{aligned} q_{j,i} &= \lambda(x_i)\Phi(x_i; \mathcal{Q}(x_i) \cap (x_j - x_i - h/2, x_j - x_i + h/2]), \quad \text{for} \quad j \neq 1, i, N_s \\ q_{1,i} &= \lambda(x_i)\Phi(x_i; \mathcal{Q}(x_1) \cap (-\infty, x_1 + h/2]) \\ q_{1,N_s} &= \lambda(x_i)\Phi(x_i; \mathcal{Q}(x_1) \cap (x_{N_s} - h/2, \infty)) \\ q_{i,i} &= -\sum_{j \neq i} q_{j,i}, \end{aligned} \tag{7}$$

where λ is the jump intensity that can be state depended.

[10] see also Chia & Skindilias (2011) for an adaptive algorithm
[11] formalised in Kushner & Dupuis (2001)

References

Andersen, L. & Andreasen, J. (2000), 'Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing', *Review of Derivatives Research* **4**, 231–262.

Bakshi, G., Cao, C. & Chen, Z. (1997), 'Empirical performance of alternative option pricing models', *Journal of Finance* **52**(5), 2003 – 2047.

Black, F. & Scholes, M. (1973), 'The pricing of options and corporate liabilities', *Journal of Political Economy* **81**(3), 637–654.

Chia, L. & Skindilias, K. (2011), 'Adaptive continuous time Markov chain approximation model to general jump-diffusions', *Submitted* .

Chourdakis, K. (2004), 'Non-affine option pricing', *Journal of Derivatives* **11**, 10–25.

Coleman, T., Li, Y. & Verma, A. (1999), 'Reconstructing the unknown local volatility function', *Journal of Computational Finance* **2**, 77–102.

Crepey, S. (2003), 'Calibration of the local volatility in a trinomial tree using Tikhonov reguralisation', *Inverse Problems* **19**, 91–127.

Derman, E. & Kani, I. (1994), 'Riding on a smile', *Risk* **7**(2).

Dumas, B., Fleming, J. & Whaley, R. E. (1998), 'Implied volatility functions: Empirical tests', *The Journal of Finance* **53**(6), 2059–2106.

Dupire, B. (1994), 'Pricing with a smile', *Risk January* pp. 18–20.

Hanson, F. B. & Westman, J. J. (2004), Stochastic analysis of jump-diffusions for financial log-return processes, *in* 'Lecture Notes in Control and Information Sciences', Vol. 280/04, Springer Berlin / Heidelberg.

He, C., Kennedy, J., Coleman, T., Forsyth, P., Li, Y. & Vetzal, K. (2006), 'Calibration and hedging under jump diffusion', *Review of Derivatives Research* **9**(1), 1–35.

Heston, S. L. (1993), 'A closed-form solution for options with stochastic volatility with applications to bond and currency options', *Review of Financial Studies* **6**(2), 327–43.

Jackson, N., Suli, E. & Howison, S. (1999), 'Computation of deterministic volatility surfaces', *Journal Computational Finance* **2**, 5–32.

Kushner, H. J. (1990), 'Numerical methods for stochastic control problems in continuous time', *SIAM Journal of Control and Optimisation* **28**(5), 999–1048.

Kushner, H. J. (1997), Numerical methods for stochastic control problems in finance, *in* M. Dempster & S.R.Pliska, eds, 'Mathematics of Derivatives Securities', Cambridge University Press, pp. 504–527.

Kushner, H. J. & DiMasi, G. (1978), 'Approximations for functionals and optimal control problems on jump-diffusion processes.', *J. Math. Anal. Applics.* **63**, 772–800.

Kushner, H. J. & Dupuis, P. G. (2001), *Numerical methods for stochastic control problems in continuous time*, Stochastic Modelling and Applied Probability, 2nd edn, Springer-Verlag, London, UK.

Lagnado, R. & Osher, S. (1997*a*), 'Reconciling differences', *Risk* **10**(4), 79–83.

Lagnado, R. & Osher, S. (1997*b*), 'A technique for calibrating derivative security pricing models: numerical solution of an inverse problem', *Journal of Computational Finance* **1**, 1325.

Lin, Y.-N., Strong, N. & Xu, X. (2001), 'Pricing FTSE-100 index options under stochastic volatility', *Journal of Futures Markets* **21**(3), 1325.

Merton, R. C. (1976), 'Option pricing when underlying stock returns are discontinuous', *Journal of Financial Economics* **3**(1-2), 125–144.

Monoyios, M. (2004), 'Efficient option pricing with transaction costs using a Markov chain approximation', *Journal of Economic Dynamics Control* **28**, 889–913.

Piccioni, M. (1987), Convergence of implicit discretisation schemes for linear differential equations with applicaiton to filtering, *in* G. Prato & L. Tubaro, eds, 'Stochastic PDE's and Applications, Lecture Notes in Mathematics 1236', Springer-Verlag.

Rubinstein, M. (1994), 'Implied binomial trees', *The Journal of Finance* **49**(3), 771–818.

Wahba, G. (1990), 'Spline models for observational data', *Soc. Industrial and Applied Mathematics* .

Zhang, X. L. (1997), 'Numerical analysis of american option pricing in a jump-diffusion model', *Mathe Oper. Res.* **22**(3), 668–690.