March 1999

PRICING EXOTIC AMERICAN OPTIONS FITTING THE VOLATILITY SMILE

M.A.H. DEMPSTER AND D.G. RICHARDS
Centre for Financial Research
Judge institute of Management Studies
University of Cambridge
Email: m.dempster@jims.cam.ac.uk & d.richards@jims.cam.ac.uk
Web: www-cfr.jims.cam.ac.uk

This paper is the latest in a series applying a new theoretical and computational method for American option valuation based on linear programming (LP). Earlier papers have treated the analytic and computational foundations, application to fast American put option valuation and the development of structured LP solution techniques for very fast valuation of (path-dependent) exotic American options such as lookbacks and Asians. In this paper we treat theoretically and numerically the inverse problem of determining the local underlying volatility from vanilla option prices and we use this in conjunction with the fast LP solver to value illustrative exotic options – fixed-strike Asians – fitting the current volatility (smile) surface implied by the market (for S&P 500 and FTSE 100 index options). Ongoing research involves similar valuation of lookbacks and barrier options fitting the smile.

1 Introduction

This paper continues a line of research concerned with the fast valuation of American options begun in [18, 34] and continued in [19, 20, 44] (see also [47]). These *linear programming* (LP) methods achieve valuation of vanilla American put options at tree speed – but with much higher accuracy [20] – through the use of a special version of the revised simplex method. This special algorithm makes use of the tridiagonal structure of the finite-difference discretization of the Black-Scholes operator, a novel basis factorization and the nature of the optimal exercise boundary to create a pricing algorithm which is essentially *linear* in the number of discretization steps in space or time with the other held fixed.

These matters are reviewed in §2 of the paper, which treats theory, basic numerical methods and our variable coefficient tridiagonal simplex algorithm in some detail, together with the degenerate PDE approach [5, 52] to valuing market-traded discretely sampled exotic options. Section 3 discusses both theory and numerical methods for fitting option values to the local volatility surface implied by option values in the market. This is an area pioneered in [21, 29, 45] and – although some new theoretical proofs are given and reliable numerical methods developed in the paper – its treatment is mainly seen here as a vehicle to demonstrate the generality and efficiency of the LP valuation algorithm in §2. In §4 results for S&P 500 and FTSE 100 exotic American index options – fixed-strike Asian puts – are presented to substantiate these claims. First, results for vanilla European and American options fitting the smile are presented in order to evaluate potential pricing errors in fitting the local volatility surface. Conclusions are drawn in §5 and directions of current and future work indicated.

2 LP Valuation of American and Discretely Sampled Exotic Options

First we review briefly the formulation of the American put option valuation problem presented in [17, 19, 20, 34]. The problem is a classical optimal stopping problem which may be formulated as a free-boundary problem by considering the domain properties of the problem. Removing any explicit reference to the free boundary, the option value may be seen to be the unique solution of an order complementarity problem by considering its equivalent formulation as a variational inequality and utilising standard results for coercive operators. Finally, the value is the solution of an abstract linear programme which can be solved with standard LP techniques upon suitable domain truncation and discretization.

2.1 Theory

Consider the standard Black-Scholes [6] economy, where we have two financial instruments – a 'risky' asset with *price* $\mathbf{S}$ modelled by a *geometric Brownian motion* (GBM) and a savings account whose balance is continuously compounded at a constant *risk-free rate* $r \geq 0$.

An *equivalent martingale* probability *measure* (EMM) $\mathbb{Q}$ (see Harrison-Kreps-Pliska [30, 31]) may be defined under which the discounted stock price process $e^{-rt}\mathbf{S}(t)$ is a *martingale* and the *stochastic differential equation* (SDE) for the stock price process becomes the GBM

$$\frac{d\mathbf{S}}{\mathbf{S}} = rdt + \sigma d\tilde{\mathbf{W}} \quad t \in [0, T] \quad S(0) > 0, \tag{1}$$

where $\sigma > 0$ is the constant *volatility* of the stock price and $\tilde{\mathbf{W}}$ is a Wiener process under $\mathbb{Q}$, which is also known as the *risk-neutral* measure.

An *option* is a risky asset whose value is determined entirely by other underlying risky assets and hence is a *derivative security*. A *European* (*vanilla*) *call* or *put option* confers the right (but not the obligation) to the holder to buy or sell respectively one unit of the asset for a price K, the *strike price*, only at a *maturity* date T. The *American* equivalent on the other hand may be exercised at any *exercise time* $\tau \in [0, T]$. Since under our assumptions an American call stock option will be optimally held to maturity, we concentrate on obtaining a formulation of the *American put* problem which is suitable for numerical solution. We define the *value function* $v : \mathbb{R}^+ \times [0, T] \to \mathbb{R}$, giving an option's fair *value* $v(x, t)$ to the holder at stock price $x > 0$ and time $t \in [0, T]$. This value is partially determined by the *payoff* function $\psi : \mathbb{R}^+ \to \mathbb{R}$, which for the American put is defined to be $\psi(\mathbf{S}(\boldsymbol{\tau})) := (K - \mathbf{S}(\boldsymbol{\tau}))^+$ and is received by the holder upon exercise at a general *stopping time* $\boldsymbol{\tau} \in [0, T]$.

The value function of an American put option can be formulated as the solution of a classical *optimal stopping problem* – choose the stopping time $\boldsymbol{\rho}(t)$ which maximises the conditional expectation of the discounted payoff – and may be shown to be the first time the value falls to the payoff at exercise, *viz.*

$$\boldsymbol{\rho}(t) := \inf \{ s \in [t, T] : \ v(\mathbf{S}(s), s) = \psi(\mathbf{S}(s)) \} . \tag{2}$$

The domain of the value function can thus be partitioned into a *continuation region* $\mathcal{C}$, on which the option has value greater than the payoff for early exercise, and a *stopping region*

$\mathcal{S}$, where the value equals the payoff since exercise occurs at the first time that the value falls to the payoff. Hence

$$\mathcal{C} := \big\{(x,t) \in \mathbb{R}^+ \times [0,T): \quad v(x,t) > \psi(x)\big\} \tag{3}$$

and

$$\mathcal{S} := \big\{(x,t) \in \mathbb{R}^+ \times [0,T): \quad v(x,t) = \psi(x)\big\}\,. \tag{4}$$

On the continuation region, the value function satisfies the Black-Scholes PDE

$$\mathcal{L}_{BS}v + \frac{\partial v}{\partial t} = 0 \tag{5}$$

for $(x,t) \in \mathbb{R}^+ \times [0,T]$, where $\mathcal{L}_{BS} := \frac{1}{2}\sigma^2 x^2 \frac{\partial^2}{\partial x^2} + rx\frac{\partial}{\partial x} - r$, since the discounted stopped price process of the option is a martingale, whilst as soon as the process crosses into $\mathcal{S}$, $v = \psi$ and

$$\mathcal{L}_{BS}v + \frac{\partial v}{\partial t} \leq 0 \tag{6}$$

to preclude arbitrage. Hence if we require the opposite inequality to (6) we have

$$\left(\mathcal{L}_{BS}v + \frac{\partial v}{\partial t}\right) \wedge (v - \psi) = 0 \tag{7}$$

on the whole domain $\mathbb{R}^+ \times [0,T)$, where $\wedge$ denotes the pointwise minimum of two functions. We now have a *free-boundary* formulation where $v(x,t) = \psi(x,t)$ for (x,t) on the optimal stopping or *exercise boundary.* We can thus remove any reference to the optimal stopping boundary by formulating the problem in terms of (7) as a *linear order complementarity problem* (OCP), using the log-transformed stock price variable $\xi := \ln x$, with respect to which the Black-Scholes operator is given by $\mathcal{L}v + \frac{\partial v}{\partial t}$, where $\mathcal{L}$ is the constant coefficient elliptic operator

$$\mathcal{L} := \frac{1}{2}\sigma^2\frac{\partial^2}{\partial \xi^2} + \left(r - \frac{1}{2}\sigma^2\right)\frac{\partial}{\partial \xi} - r \tag{8}$$

and v is now the option value as a function of ξ. The various inequalities carry through the domain transformation and the new payoff function is given by $\psi(\xi) := \big(K - e^{\xi}\big)$. As shown in [19] the American put value function is the unique solution to

$$(\text{OCP}) \begin{cases} v(\cdot,T) = \psi \\ v \geq \psi \\ \mathcal{L}v + \frac{\partial v}{\partial t} \geq 0 \\ \big(\mathcal{L}v + \frac{\partial v}{\partial t}\big) \wedge (v - \psi) = 0 \quad \text{a.e. in } \mathbb{R} \times [0,T] \end{cases} \tag{9}$$

posed in a suitable *vector lattice* Hilbert space, which is a Hilbert space H with inner product $\langle\cdot,\cdot\rangle$ and partial order defined by a positive cone P such that for any points x and y the maximum $x \vee y$ and the minimum $x \wedge y$ exist in the given order [9, 15]. Dempster and Hutton [19] (see also [36]) use another equivalent formulation of the value function problem as a *variational inequality* (VI) to show the uniqueness of the solution to (OCP) if the differential operator is *coercive*, i.e. $\exists \alpha \in \mathbb{R}^+$ s.t. $\langle u, \mathcal{L}u \rangle \geq \alpha \|u\|^2 \quad \forall u \in H$. They show that the value function, as the unique solution to (OCP), can be expressed as the unique solution of an abstract *linear programme* given by

$$\text{(LP)} \qquad \inf_v \langle v, c \rangle \text{ s.t. } v \in \mathcal{F} \text{ for any } c > 0 \text{ a.e. on } \mathbb{R} \times [0, T], \tag{10}$$

where

$$\mathcal{F} := \left\{ v : v(\cdot, T) = \psi,\ v \geq \psi,\ \mathcal{L}v + \frac{\partial v}{\partial t} \geq 0 \right\} \tag{11}$$

since the linear operator $\mathcal{L}$ on the Hilbert space H is of *type-Z*, i.e. $\langle v, y \rangle = 0 \ \Rightarrow\ \langle v, \mathcal{L}y \rangle \leq 0 \ \forall v, y \in H$.

From this abstract LP formulation the problem can be reduced from infinite to finite dimensions through space and time discretizations and the resulting ordinary LP can be solved to find a numerical approximation to the value function. To this end, the value function on $\mathbb{R} \times [0, T]$ is restricted to a finite region $[L, U] \times [0, T]$ with explicit conditions on the boundaries of the domain. Then, defining a localised inner product with integration over the reduced domain, we have a *localised* abstract LP with new constraint set

$$\mathcal{F} := \left\{ v : v(L, \cdot) = \psi(L), v(U, \cdot) = \psi(U), v(\cdot, T) = \psi,\ v \geq \psi,\ \mathcal{L}v + \frac{\partial v}{\partial t} \geq 0 \right\} \tag{12}$$

which in the limit, as $L \to -\infty$ and $U \to \infty$, converges to the solution of the abstract problem [19].

2.2 Numerical Methods

We approximate the value function by a function which is piecewise constant on rectangular intervals between points in a regular lattice of dimension $I \times M$. Denote the value at a general point $(L + i\Delta\xi, T - m\Delta t)$ by $v_i^m := v\,(L + i\Delta\xi, T - m\Delta t)$, where $m \in \{0, 1, \ldots, M\} =: \mathcal{M}$ and $i \in \{0, 1, \ldots, I\} =: \mathcal{I}$. Approximating the partial derivatives by standard *Crank-Nicolson finite differences* [52] we obtain a discrete form of (OCP) which, upon collapsing the space index, can be rewritten in matrix form. The complementarity condition (line 3 of (9)) is given in matrix form by

$$\left(Bv^{m-1} + Av^m - \phi\right) \wedge (v^m - \psi) = 0 \qquad m \in M \backslash \{0\}, \tag{13}$$

where A and B are $I-1$ square tridiagonal matrices with constant nonzero entries denoted by $\{a, b, c\}$ and $\{d, e, f\}$ respectively, and

$$v^m := \begin{pmatrix} v_1^m \\ \vdots \\ v_{I-1}^m \end{pmatrix} \quad \psi := \begin{pmatrix} \psi_1^m \\ \vdots \\ \psi_{I-1}^m \end{pmatrix} \quad \phi := \begin{pmatrix} -(a+d)\psi_0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}. \tag{14}$$

It is easy to see that the matrix A represents the discrete approximation of the continuous linear type-Z operator $\mathcal{L}$, so it is necessary to find conditions for the matrix A to be type-Z. By definition [9] a matrix is *type-Z* if it has non-negative off-diagonal elements, which in the case of A occurs when $|r - \sigma^2/2| \leq \sigma^2/\Delta\xi$ and can be satisfied by adjusting the number of space steps I in the discretization. From this condition it can also be shown that A is coercive [19, 36]. Hence we can formulate the discretized OCP by considering the finite (time step) sequence of order complementarity problems

$$\begin{aligned} &v^m \geq \psi \\ &Bv^{m-1} + Av^m - \phi \geq 0 \\ &\left(Bv^{m-1} + Av^m - \phi\right) \wedge (v^m - \psi) = 0 \end{aligned} \tag{15}$$

with equivalent sequence of ordinary LPs

$$\begin{aligned} &\min c' v^m \\ &\text{s.t.} \quad v^m \geq \psi \\ &Av^m \geq \phi - Bv^{m-1} \qquad m = 1, \dots, M. \end{aligned} \tag{16}$$

The LP formulation can be solved either directly or iteratively and the interested reader can find comparisons of solution methods in [19, 34].

We describe next a simplified revised simplex method for solution of the LP formulation of the vanilla American put option valuation problem which takes advantage of the tridiagonal structure of the constraint matrix of (16), formed from standard Crank-Nicolson finite difference approximations, to produce a fast accurate direct solution method. For more details on the terminology in this section see a standard LP text such as [41].

To rewrite (16) in *standard form* we define a new variable u^m which is the value of the option in excess of the payoff function, $u^m := v^m - \psi$. Substituting gives

$$\begin{array}{lll} \min & c'u^m & \\ \text{s.t.} & u^m & \geq 0 \\ & Au^m & \geq b, \end{array} \tag{17}$$

where the right-hand side vector b is given by $b := \phi - B\left(u^{m-1} + \psi\right) - A\psi$.

Setting $n := I - 1$, we convert (17) to an underdetermined $n \times 2n$ system of linear equations by adding non-negative *slack* variables $s := (s_1, s_2, \dots, s_n)$, giving

$$\min\left(c' \ 0'\right)\begin{pmatrix} u^m \\ s \end{pmatrix} \quad \text{s.t.} \ (A \quad -I)\begin{pmatrix} u^m \\ s \end{pmatrix} = b, \quad u^m \geq 0, \ s \geq 0. \tag{18}$$

The constraints of (18) describe a polytope in $\mathbb{R}^{2n}$, with the (unique) optimal solution of (18) at a vertex of this polytope. A vertex may be identified by setting n of the (slack and real) variables (*non-basics*) to zero and solving the modified system $D\bar{u} = b$ for the remaining n *basic* variables, where D is the $n \times n$ *basis* matrix constructed from the columns of the constraint matrix corresponding to the basic variables and $\bar{u}$ is the corresponding vector of basic slack and real variables.

We first choose an initial basis, which simply amounts to excluding *nb real* (i.e. not slack) non-basic variables from the basis so that it comprises $\bar{u}_{nb} = \left(s_1 \dots s_{nb} \ u^m_{nb+1} \dots u^m_n\right)'$. Note that we are assuming the connectedness of the index sets of the real basic variables and the slack basic variables as subsets of $\mathbb{N}$. This is implied by the connectedness of the stopping and continuation regions in $[L, U] \times [0, T]$ (see Figure 1). We also assume that the optimal basis contains u^m_n, which can be guaranteed by appropriate indexing, given connectedness. With this basis specified, we next find the solution of the linear system

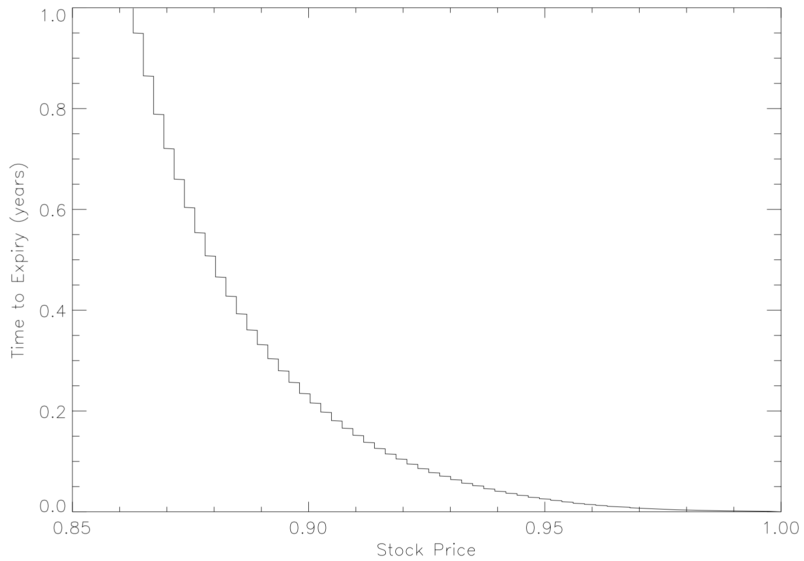

Figure 1: The calculated optimal exercise boundary for the American put option $r := 10\%$, $\sigma := 20\%$ with discretization $M := 1000$ and $I := 1200$ and stock price range $[0.37, 7.39]$.

$$D\bar{u}_{nb} = b,$$

where

$$D := \begin{pmatrix} -1 & & & & & \\ & \ddots & & & & \\ & & -1 & a & & \\ & & & b & \ddots & \\ & & & c & \ddots & a \\ & & & & \ddots & b & a \\ & & & & & c & b \end{pmatrix}. \tag{19}$$

This solution may be found as

$$\bar{u}_{nb} = \begin{pmatrix} -I_{nb-1} & \\ & \tilde{D}^{-1}_{n-nb+1} \end{pmatrix} b, \tag{20}$$

where I_{nb-1} is the $nb-1$ identity matrix and $\tilde{D}^{-1}_{n-nb+1}$ is the tridiagonal $n-nb+1$ square matrix given by

$$\tilde{D}_{n-nb+1} := \begin{pmatrix} -1 & a & & & \\ & b & \ddots & & \\ & c & \ddots & a & \\ & & \ddots & b & a \\ & & & c & b \end{pmatrix}. \tag{21}$$

Since the objective function coefficients $c_1, \dots, c_n$ are arbitrary positive numbers, we can choose these so that a basis with $n - nb$ real variables always produces a solution (20) with a smaller optimal value than a basis with $n - nb - 1$ real variables. Therefore, from a given feasible basis one may find an optimal basis by repeatedly adding the next real variable u^m_{nb} into the current basis until the corresponding solution (20) becomes infeasible – then the new optimal basis is that with one less real basic variable than the first infeasible basis. We may start the iterative process for u^m from the previous time-step's solution vector u^{m-1} – in the case of the American put we know that this will be feasible since the exercise boundary has a convex graph in $[L, U] \times [0, T]$ (see Figure 1). With options for which this graph has a positive slope, one can reverse the iterative procedure by removing real variables from the basis until feasibility is achieved.

The above procedure requires repeated solution of the tridiagonal system (20) for a sequence of basic real variables. This may be done efficiently by factorization of the tridiagonal matrix $\tilde{D} := LU$ defined by (21) by noting that the factorization involving lower and upper triangular bidiagonal matrices L and U respectively need only be computed once for $nb = 1$, i.e. $\tilde{D}_n := A$ [20]. We need only compute the factorization for the matrix $\tilde{D}_n := A$ of size n and read-off the required factorizations of smaller matrices $\tilde{D}_{n-nb+1}$ as we progressively increase the number of real variables $n - nb$ in the basis. This algorithm can be implemented efficiently in any computing system using only 3 storage arrays containing the nontrivial coefficients of the LU decomposition along the 3 diagonals.

This procedure is suitable for any standard constant parameter Black-Scholes type formulation, but we now outline a procedure which yields significant computational savings for valuation problems with volatility and drift parameters which are functions of time. It also incorporates a technique for the solution of problems with non-constant constraint matrix coefficients such as those involving the untransformed Black-Scholes PDE, which has coefficients given by functions of the underlying asset price, or for exotic option pricing problems, where the coefficients vary with the third variable representing the path-dependency. In [20] results are presented for this updating procedure which show that even for a general constraint matrix the procedure out-performs standard commercial LP solvers by orders of magnitude.

The greatest efficiency saving in the standard LU factorization follows from the observation that for the constant coefficient constraint matrix the factorization need only be performed once at the outset of the algorithm. This would not be the case using the above technique with time-dependent parameters, for these would require a full factorization of the initial basis at each time-step. With the LU formulation it is not so easy a task to update the factorization with the introduction of new real basic variables due to the recursive 'above-diagonal' nature of the computation of the diagonal of the U matrix. We therefore 'reverse' the factorization to allow for computationally efficient updating. Define

$$\bar{D}_{nb,n} := \begin{pmatrix} -1 & a_{nb} & & & & \\ & b_{nb+1} & \ddots & & & \\ & c_{nb+2} & \ddots & \ddots & & \\ & & \ddots & \ddots & a_{n-2} & \\ & & & c_{n-1} & b_{n-1} & a_{n-1} \\ & & & & c_n & b_n \end{pmatrix} \tag{22}$$

to be the $n - nb + 1$ square submatrix of the basis matrix corresponding to (20) when nb non-basic real variables have been excluded. The subscripts in (22) represent entries of a general $n - nb + 1$ square tridiagonal matrix with entries which vary with their indices, for example, to be dependent on the asset value. We factorize (22) by writing $\bar{D}_{nb,n} = U_{nb,n} L_{nb,n}$ where $U_{nb,n}$ and $L_{nb,n}$ are upper and lower triangular bidiagonal $n - nb + 1$ square matrices respectively, with $U_{i\ i} := 1 \quad i = nb, \ldots, n$. With this 'reversed' factorization we remove the need to recursively calculate all the factor matrices upon the introduction of a new real variable - instead we perform a simple update.

Setting c_{nb} and b_{nb} equal to zero for notational simplicity, the factorization proceeds as follows. At each iteration we start from a basis with $n - nb$ real variables and factorize by backwards recursion as

$$\begin{aligned} U_{n\ n} &= 1 \\ L_{n\ n-1} &= c_n \\ L_{n\ n} &= b_n \\ \vdots\ \ &\ \ \vdots \\ L_{i\ i-1} &= c_i \\ U_{i\ i+1} &= \frac{a_i}{L_{i+1\ i+1}} \\ L_{i\ i} &= b_i - U_{i\ i+1} L_{i+1\ i} & \qquad i = n-1, \ldots, nb+1 \\ U_{i\ i} &= 1 \\ \vdots\ \ &\ \ \vdots \\ U_{nb\ nb+1} &= \frac{a_{nb}}{L_{nb+1\ nb+1}} \\ L_{nb\ nb} &= -1. \end{aligned} \tag{23}$$

When another real variable enters the basis, we perform a simple update by increasing by one the dimension of the square matrices, calculating the new columns of L and U corresponding to the new variable, and re-calculating certain elements in the previous columns, *viz.*

$$
\begin{aligned}
L_{nb+1\ nb} &= c_{nb+1} \\
L_{nb\ nb-1} &= 0 \\
L_{nb\ nb} &= b_{nb} - U_{nb\ nb+1} L_{nb+1\ nb} \\
\vdots \quad & \quad \vdots \\
U_{nb-1\ nb} &= \frac{a_{nb-1}}{L_{nb\ nb}} \\
L_{nb-1\ nb-1} &= -1.
\end{aligned}
\tag{24}
$$

The number nb of real non-basic variables is then decremented and the procedure continues as above. The full UL factorization has the same computational complexity as the LU decomposition for a full factorization, but only three floating point operations are required at each update using (24).

To gain an understanding of the exact computational savings of the above methods, first consider the complexity of the one-factor American put option valuation problem after transformation to the constant-coefficient Black-Scholes operator. At each time-step the maximum number of real variables which can enter the basis is given by $\lfloor \frac{\ln K - L}{\Delta x} \rfloor$, where K is the strike price, L and Δx are respectively the lower bound and space step size of the discretization of the space domain and $\lfloor \cdot \rfloor$ denotes integer part. Thus we have $\mathcal{O}(I)$ possible new basis variables, i.e. iterations, at each time step, where I is the number of points in the spatial discretization. In fact, after the first few time steps – where the exercise boundary has greatest curvature away from $\ln K$ (see Figure 1) – at most one new basic variable enters at each time step. Far from maturity, calculations for several time steps may utilize the *same* basis. Each iteration requires $\mathcal{O}(n)$ operations to solve, where $n \leq I$, giving $\mathcal{O}(I)$ operations at each time step. Hence the space complexity of the algorithm is *linear* and the total operation count is $\mathcal{O}(TI)$, where T is the number of time-steps.

For the updating technique the calculations result in a similar complexity, but extra solution time is needed for the dynamic allocation of the UL factorization at each iteration. For the full recalculation method it is necessary to include the UL factorization calculation at each iteration, resulting in an extra $\mathcal{O}(I)$ operations at each time point but still $\mathcal{O}(I)$ complexity – a significant saving over the $\mathcal{O}(I^3)$ operations required for a full $I \times I$ matrix LU factorization and equation solution.

Results for the constant coefficient method and for the non-constant coefficient updating technique are reported in [20], along with results for a complete calculation of the full LU factorization at each iteration to highlight the overheads of using general commercial solvers.

2.3 Discretely Sampled Exotic Options

An *exotic option* is any derivative security which has a path-dependent component in its payoff at exercise. *Vanilla options* on the other hand have payoffs which are at most functions of the stock price at exercise time. We may also formulate exotic option valuation problems as linear programs, with new values dependent on the underlying stock price, time and an additional 'independent' variable which encapsulates the required path information.

This PDE is derived by augmenting the state-space with a new independent variable representing the path-dependent quantity to create a degenerate two-dimensional PDE [5] with no diffusion (i.e. second derivative) term in the new variable. It can be shown (see [52]) that when a path-dependent quantity is sampled discretely on a finite number of occasions (as for traded options) the option value satisfies a fixed parameter one-dimensional Black-Scholes equation with jump conditions across sampling dates. As a result the degeneracy can be removed and we can express discretely-sampled exotics in LP form. The problem must still be solved in two space dimensions, but the extra variable enters only as a parameter in the valuation problem.

We outline here the formulation of a general American exotic option in a discretely-sampled setting using a dynamic programming algorithm for the option valuation based on the unifying framework of [52]. Denote by $V(S, M, t)$ the *value function* of the option with $V : \mathbb{R}^+ \times \mathbb{R}^+ \times [0, T] \to \mathbb{R}$, where S denotes the asset price and M denotes the current value of the additional path-dependent variable. We assume that the asset price is sampled on N occasions during the life of the option with maturity T. Denote by M_n the *observed asset price* at the *sampling date* t_n, $n = 0, \ldots, N-1$. For completeness define $t_N := T$ and assume that the sampling begins at time 0,[1] so that $t_0 = 0$ and $M_0 = S(0)$. The variable M_n is constant throughout the period $[t_n, t_{n+1})$, since no sampling takes place until time t_{n+1}. Effectively M_n is simply a parameter in the formulation during this period, and any randomness in the model is due to the asset price process. The Black-Scholes PDE will thus be satisfied within the period with jump conditions applied at sampling dates, see [25, 52] for more details.

Across a general sampling date t_n the path-dependent variable is updated from a value M_{n-1} just prior to the date to a value M_n at the sample date. To avoid arbitrage opportunities the option value must be continuous across sampling dates for any particular realization of the asset. This leads to the jump condition

$$V\left(S, M_{n-1}, t_n^-\right) = V\left(S, M_n, t_n\right) \qquad n = 1, \ldots, N-1, \tag{25}$$

where t_n^- and t_n are times immediately before and at the sampling date t_n. In the time interval $[t_n, t_{n+1})$ to the next sampling date, V satisfies the *augmented Black-Scholes PDE* given by

$$\frac{\partial V}{\partial t} + \sum_{i=1}^{N} \delta(t - t_i) f(S) \frac{\partial V}{\partial M} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0, \tag{26}$$

where $\delta(\cdot)$ denotes the *Dirac delta function* and $f(S)$ is a function to be determined for each specific exotic option.

We consider the final period $[t_{N-1}, T]$ and use a dynamic programming algorithm to determine values for earlier periods. As in the American put case, but with increased dimension, the American exotic valuation domain $\mathbb{R}^+ \times \mathbb{R}^+ \times [t_{N-1}, T]$ can be partitioned into a continuation region $\mathcal{C}_N$ and a stopping region $\mathcal{S}_N$ and we can establish the existence of an optimal exercise boundary. In this period we must have from arbitrage considerations

$$V(S, M_{N-1}, t) \geq \psi(S, M_{N-1}) \qquad t \in [t_{N-1}, T], \tag{27}$$

for any possible value of M_{N-1} with V and $\frac{\partial V}{\partial S}$ continuous. The boundary at $S = 0$ is an absorbing boundary since the asset price follows GBM and if the asset has zero value it will remain zero. If the option is held until maturity in this case, then the value at exercise is

[1] The implementation of any sampling scheme is computationally straightforward, so that very general schemes can be solved in this manner.

equal to the payoff and so at a time $t \in [t_{N-1}, T]$ the option value is given by the discounted payoff

$$V(0, M_{N-1}, t) = e^{-r(T-t)}\psi(0, M_{N-1}) \qquad t \in [t_{N-1}, T]. \tag{28}$$

This contradicts (27) and so the option must be *stopped*, i.e optimally exercised, when the asset price reaches 0.

To complete the formulation of the discretely-sampled exotic option value in the final period we require a terminal condition and boundary conditions at $S = 0$ and as $S \to \infty$. In the final period $[t_{N-1}, t_N]$ our terminal condition is that the value of the option equals the payoff at maturity.

The boundary condition at $S = 0$ is given by (28). As $S \to \infty$ the value of the option tends to zero monotonically, since at maturity the option value is zero if $S \geq M_{N-1}$. It is sufficient for this formulation to say that the option value can grow at most linearly[2] with S as $S \to \infty$. Hence we implement the boundary condition

$$\frac{\partial^2 V}{\partial S^2} \to 0 \text{ as } S \to \infty. \tag{29}$$

Again we log-transform the primitive variables ($\xi := \ln S, \zeta_{N-1} := \ln M_{N-1}$) and formulate the valuation problem with fixed ζ_{N-1} as an OCP with respect to the transformed operator $\mathcal{L} := \frac{1}{2}\sigma^2 \frac{\partial^2}{\partial \xi^2} + \left(r - \frac{1}{2}\sigma^2\right) \frac{\partial}{\partial \xi} - r$, defining a new partition with regions $\tilde{\mathcal{C}}_N$ and $\tilde{\mathcal{S}}_N$. Thus the American exotic valuation problem in the final period may be formulated in terms of the transformed value function $V := V\left(e^{\xi}, e^{\zeta_{N-1}}, t\right)$ as the unique solution of the order complementarity problem

$$\text{(OCP)} \begin{cases} V(\cdot, \cdot, T) = \tilde{\psi} \\ V \geq \tilde{\psi} \\ \mathcal{L}V + \frac{\partial V}{\partial t} \geq 0 \\ \left(\mathcal{L}V + \frac{\partial V}{\partial t}\right) \wedge (V - \tilde{\psi}) = 0 \quad \text{a.e. in } \mathbb{R} \times \mathbb{R} \times (t_{N-1}, T], \end{cases} \tag{30}$$

where $\tilde{\psi}(\xi, \zeta_{N-1})$ is the payoff function and V now denotes the option value as a function of ζ_{N-1} and ξ. This puts us in a framework equivalent to the vanilla American put in Section 2, but with the additional parameter ζ_{N-1}, and hence we can show equivalence to an abstract LP for each value of $\zeta_{N-1} \in (-\infty, \infty)$. The problem must now be solved for all possible values of the parameter ζ_{N-1}. Applying the jump conditions (25) at t_{N-1} to obtain the terminal value $V(S, M_{N-2}, t_{N-1}^-)$, the argument may be repeated for the period $[t_{N-2}, t_{N-1}]$ and, by backwards recursion, eventually for the period $[0, t_1]$.

3 Fitting the Volatility Smile

In §2.1 we described the Black-Scholes model which is an idealised pricing environment. Black-Scholes theory assumes that stock prices are lognormally distributed and stochastic, yet the future volatilities and interest-rates are deterministic. However in recent years the market has been pricing options violating the assumption of known constant volatilities, and the Black-Scholes model has been adapted to deal with these deviations from the lognormal

[2] For further discussion see [52] (pp. 212–214).

world. In this section we suggest alternate ways of looking at the 'real-world' by relaxing some of the assumptions of Black-Scholes theory and calibrating the option pricing algorithm to market data.

3.1 Empirical Evidence for Non-Constant Volatility

The Black-Scholes economy has one unobservable quantity - the *volatility* parameter σ. To be able to price derivatives some suitable value of this parameter must be inferred. Particularly since the stock market crash of 1987 the volatility of equity options has exhibited variation both with the strike price and the options' maturity. The time dependence gives rise to a *term structure* of volatility and the curvature with respect to the strike price is termed the *volatility smile* (or *skew*). Both of these effects highlight the markets' deviation from the assumption in the Black-Scholes economy that the future asset price has a constant variance lognormal conditional probability density. Since 'zero-probability' events such as the stock market crash of 1987, or the market turmoil in the Far East in 1998, are so recent in the memory of market-makers, it is difficult to reconcile these events with a short-tailed distribution. In effect the market is grafting long tails onto the probability density and changing its temporal shape to cope with these memories.

Several approaches have been suggested in the literature to model this behaviour. The first approach is to treat the volatility as an *additional stochastic variable* with the aim of recovering the volatility parameter from the model. This approach was first suggested by Hull and White [32] and is also explored in [12, 4]. Whilst giving the model the ability to deal with non-constant effects, this approach is difficult to fit to the data and is not arbitrage-free. It also introduces an additional dimension to the pricing problem with the obvious additional computational complexity that this entails.

The second approach is to allow the volatility $\sigma := \sigma(S, t)$ to be a variable which is both state and time dependent. Early methods specified a functional form for the volatility but did not generally fit the market data sufficiently. However, by starting from the market data and backing-out the local volatilities which are consistent with the market, this model can be made to price the market exactly. This is commonly termed an *inverse problem*[39, 38]. The most popular structures on which this local volatility is determined are *binomial* or *trinomial* trees, which allow specification of nodal transitional probabilities to fit the smile. These *implied tree* approaches differ in complexity and accuracy of fit, with some fitting a single terminal maturity probability distribution [48, 45, 46] and others attempting a multiple-maturity fit. Dupire [28] introduced a continuous-time theory involving the adjoint PDE to the Black-Scholes equation [40] which we summarise below. This work has been extended by Derman & Kani [22, 23] whose latter paper also attempts to model the volatility surface through calendar time using *shocks* - in essence a combination of an implied tree and a stochastic volatility model.

The methods used to fit the market data to a tree or lattice are all prone to the same instabilities. Generally the data can imply unreasonable (e.g. negative or large) values of the local volatility, which may create negative transitional probabilities which necessarily allow arbitrage possibilities in the model. Several methods have been suggested to calibrate the data to models. *Filtering* is one such method, but so far in practice filtering has only been done in an *ad hoc* manner with any unstable volatility values (or equivalently negative transitional probabilities) simply being ignored or set to zero. For examples of methods in this category see [49, 50] which apply principal component analysis to determine the number of factors in the Derman and Kani [23, 24] model.

The general inverse problem is *ill-posed* since the number of volatility parameters to be found far outnumbers the limited number of available option prices in the market. It is often assumed that a continuum of European call option prices $C(K,T)$ are available for all strikes and maturities, although in practice this is not the case, but the assumption is usually justified by the use of interpolation and extrapolation of given market data to obtain the call option prices for any required maturity and strike price.

Recent papers have implemented *regularization* methods [38] to make the inverse problem stable. The method of *Tikhonov regularization* [8, 7, 40] uses an optimization approach to add a smoothing measure which ensures that the inverse problem is well-posed and that the volatility surface solution is the unique minimizer of some goodness-of-fit measure relative to observed option prices. Another such approach was introduced in [13] whose authors apply regularization to a function interpolating and extrapolating the observed market data. Other regularization approaches have been taken in [39, 35] - the latter uses *maximum entropy regularization* as the stabilizing functional based on a prior estimate of the local volatility.

Finally, the recent article by Andersen and Brotherton-Ratcliffe [2] should be mentioned because it is closest in spirit to the approach utilised here. Their method involves forward induction [37] of Arrow-Debreu prices to generate the *local volatility* – similar to the approach of Dupire [29] – through an implicit finite difference approach. The continuous coefficients in the Black-Scholes PDE are replaced by discrete equivalents derived from the market data, specifically call option prices, bond prices and forward prices. Since the same market data is used in this approach as in the implied tree approach the same data discrepancies occur.

Thus several approaches have been pursued to model the smile effect, but all suffer from the consequences of inconsistent data and will not price *all* options correctly in the face of these data problems. For some methods it is enough to price a subset of options accurately, but even in these cases there is no guarantee that arbitrage opportunities will not occur. The modelling approach used in this paper is not claimed to be the most accurate or efficient, but is simply used to highlight how versatile the LP pricing method is in the face of a degenerate, ill-posed problem with non-constant coefficients.

3.2 Continuous-Time Volatility Theory

An arbitrage-free local volatility surface in *continuous time* can be inferred from market data - in particular the prices of European call options. This theory was first derived by [29, 28] and has since been given more formal treatments in [22, 21, 23]. The main idea is that there exists an *adjoint* [40] or *dual* PDE to the Black-Scholes PDE which has the strike price K and maturity T as the independent variables, instead of the asset price S and current time t. The PDE can be derived through consideration of the conditional probability distribution of the underlying stochastic process and the forward PDE satisfied by this probability density.

As discussed above we assume that a *continuum* of European call option prices for all strikes and maturities are available from the market, i.e. there exists $C(K,T)$ for all $K,T \in \mathbb{R}^+$. In practice of course this is not the case, but any 'gaps' in the data can be filled by interpolation or extrapolation techniques and we will deal with any arbitrage violations in these approximated values later. The underlying asset price is assumed to follow a one-factor diffusion process under the risk-neutral measure $\mathbb{Q}$, but now with non-constant volatility, *viz.*

$$d\mathbf{S} = \mathbf{S}r(t)dt + \mathbf{S}\sigma(\mathbf{S},t)d\mathbf{W}, \tag{31}$$

where r is the risk-free rate assumed to be at most dependent on time and σ is now allowed to be dependent on both state variables. The price of a European call option can be written in terms of an expectation under $\mathbb{Q}$ in terms of the *conditional probability density function* $p(s,T|S,t)$ of the underlying asset $\mathbf{S}$ having value s at time T given that the asset price is S at time t. Hence

$$C(S,t;K,T) = \exp\left(-\int_t^T r(u)du\right)\int_0^\infty p(s,T|S,t)(s-K)^+ds \tag{32}$$

$$= P(t,T)\int_K^\infty p(s,T|S,t)(s-K)ds. \tag{33}$$

where $t \leq T$ and the *discount factor*

$$P(t,T) := \exp\left(-\int_t^T r(u)du\right).$$

Breeden and Litzenberger [10] showed that the European call option price and this conditional probability density were related by differentiation of the call price.

Proposition 1
The conditional probability density $p(K,T|S,t)$ is given by

$$p(K,T|S,t) = \frac{1}{P(t,T)}\frac{\partial^2 C(S,t;K,T)}{\partial K^2} \tag{34}$$

where $C(S,t;K,T)$ is the European call price given by (33).

The function $p(K,T|S,t)$ is the risk-neutral transitional probability density and is also the *Green's function* (or *fundamental solution* [27, 51]) of the Black-Scholes PDE for the European call option value. Thus it satisfies the PDE with terminal condition $(t=T)$

$$p(K,T|S,t) = \delta(S-K),$$

where $\delta(\cdot)$ is the Dirac delta function. Since C is assumed known this density function can be found from the idealised market data.

The risk-neutral conditional probability density function, $p(K,T|S,t)$ can also be shown to satisfy the *Fokker-Planck* (or *forward Kolmogorov*) PDE, through the following theorem.

Theorem 2
The conditional probability density function $p(y,\tau|x,t)$ of a general stochastic process $\mathbf{X}(t)$ where $t \geq 0$ given by

$$d\mathbf{X}(t) = \mu(x,t)dt + \sigma(x,t)d\mathbf{W}(t)$$

satisfies the Fokker-Planck or forward Kolmogorov equation

$$\frac{\partial p(y,\tau|x,t)}{\partial \tau} + \frac{\partial\left(\mu(y,\tau)p(y,\tau|x,t)\right)}{\partial y} - \frac{1}{2}\frac{\partial^2\left(\sigma^2(y,\tau)p(y,\tau|x,t)\right)}{\partial y^2} = 0, \tag{35}$$

for fixed $(x,t) \in \mathbb{R}\times\mathbb{R}^+$ with initial condition

$$p(y,t|x,t) = \delta(x-y).$$

Proof
See a standard stochastic differential equation text such as [42], or see [44]. □

Corollary 3
The transitional probability density $p(S', T|S, t)$ of the stock price process given by (31) *satisfies the PDE*

$$\frac{\partial p(S',T|S,t)}{\partial T} = \frac{1}{2}\frac{\partial^2(\sigma^2(S',T)S'^2 p(S',T|S,t)}{\partial S'^2} - \frac{\partial}{\partial S'}\left(rS'p(S',T|S,t)\right) \tag{36}$$

with initial condition $(T = t)$ $p(S', T|S, t) = \delta(S' - S)$.

Proof
The proof follows from Theorem 2 by noting that for the process (31) the drift term $\mu(y, \tau) = r(\tau)y$ and variance $\sigma^2(y, \tau) = \sigma^2(y, \tau)y^2$. Substituting in (35) gives the result, since y is simply a value of the process at time τ so can be replaced by S' at time T. □

We now state the following corollary, originally due to Dupire [29].

Corollary 4
Given that the underlying asset price process follows (31) *the price of a European call option* $C(S, t; K, T)$ *solves the partial differential equation*

$$\frac{\partial C}{\partial T} = \frac{\sigma(K,T)^2K^2}{2}\frac{\partial^2 C}{\partial K^2} - r(T)K\frac{\partial C}{\partial K} \tag{37}$$

with boundary condition $C(S, t; S, t) = 0$.

Proof
We prove the corollary by substituting for the density $p(S', \tau|S, t)$ given by (34) into the Fokker-Planck equation (36). Thus

$$\begin{aligned}
\frac{\partial p(S',\tau|S,t)}{\partial \tau} &= \frac{\partial}{\partial \tau}\left[\frac{1}{P(t,\tau)}\frac{\partial^2 C}{\partial S'^2}\right] \\
&= \frac{1}{P(t,\tau)}\frac{\partial}{\partial \tau}\left[\frac{\partial^2 C}{\partial S'^2}\right] + \frac{\partial^2 C}{\partial S'^2}\frac{\partial}{\partial \tau}\left[\frac{1}{P(t,\tau)}\right] \\
&= \frac{1}{P(t,\tau)}\left[\frac{\partial}{\partial \tau}\left(\frac{\partial^2 C}{\partial S'^2}\right) + r(\tau)\frac{\partial^2 C}{\partial S'^2}\right] \\
&= \frac{1}{P(t,\tau)}\left[\frac{\partial^2}{\partial S'^2}\left(\frac{\partial C}{\partial \tau}\right) + r(\tau)\frac{\partial^2 C}{\partial S'^2}\right]
\end{aligned} \tag{38}$$

Similarly,

$$\frac{\partial}{\partial S'}\left(rS'p(S',\tau|S,t)\right) = \frac{r}{P(t,\tau)}\frac{\partial}{\partial S'}\left(S'\frac{\partial^2 C}{\partial S'^2}\right) \tag{39}$$

and substituting back into (35) gives

$$\frac{\partial^2}{\partial S'^2}\left(\frac{\partial C}{\partial \tau}\right) + r(\tau)\frac{\partial^2 C}{\partial S'^2} + r(\tau)\frac{\partial}{\partial S'}\left(S'\frac{\partial^2 C}{\partial S'^2}\right) - \frac{1}{2}\frac{\partial^2}{\partial S'^2}\left(\sigma^2(S',\tau)S'^2\frac{\partial^2 C}{\partial S'^2}\right) = 0. \tag{40}$$

Following [22], multiplying (40) by $(S' - K)$ and integrating with respect to S' from K to ∞ gives

$$\int_K^\infty \left[\frac{\partial^2}{\partial S'^2}\left(\frac{\partial C}{\partial \tau}\right) + r(\tau)\frac{\partial^2 C}{\partial S'^2} + r(\tau)\frac{\partial}{\partial S'}\left(S'\frac{\partial^2 C}{\partial S'^2}\right) - \frac{1}{2}\frac{\partial^2}{\partial S'^2}\left(\sigma^2(S',\tau)S'^2\frac{\partial^2 C}{\partial S'^2}\right)\right](S'-K)dS' = 0. \tag{41}$$

Then, integrating by parts and considering each term of (41) in turn

$$\int_K^\infty \frac{\partial^2}{\partial S'^2}\left(\frac{\partial C}{\partial \tau}\right)(S'-K)dS' = \frac{\partial C}{\partial \tau} + \left[(S'-K)\frac{\partial}{\partial S'}\left(\frac{\partial C}{\partial \tau}\right)\right]_{S'=\infty} \tag{42a}$$

$$\int_K^\infty r(\tau)\frac{\partial^2 C}{\partial S'^2}(S'-K)dS' = r(\tau)C(S,t;K,\tau) + \left[r(S'-K)\frac{\partial C}{\partial S'}\right]_{S'=\infty} \tag{42b}$$

$$\begin{aligned}\int_K^\infty r(\tau)\frac{\partial}{\partial S'}\left(S'\frac{\partial^2 C}{\partial S'^2}\right)(S'-K)dS' &= r(\tau)K\frac{\partial C}{\partial K} - r(\tau)C(S,t;K,\tau) \\ &\quad + \left[r(\tau)(S'-K)S'\frac{\partial^2 C}{\partial S'^2} - r(\tau)S'\frac{\partial C}{\partial S'}\right]_{S'=\infty}\end{aligned} \tag{42c}$$

$$\begin{aligned}\int_K^\infty \frac{1}{2}\frac{\partial^2}{\partial S'^2}\left(\sigma^2(S',\tau)S'^2\frac{\partial^2 C}{\partial S'^2}\right)(S'-K)dS' &= \frac{1}{2}\sigma^2(K,\tau)K^2\frac{\partial^2 C}{\partial K^2} \\ &\quad + \left[\frac{1}{2}(S'-K)\frac{\partial}{\partial S'}\left(\sigma^2(S',\tau)S'^2\frac{\partial^2 C}{\partial S'^2}\right)\right]_{S'=\infty}.\end{aligned} \tag{42d}$$

The boundary terms can be shown to tend to zero as $S' \to \infty$ under suitable justifiable assumptions. Firstly, we assume that $\frac{\partial C(S,t;S',t)}{\partial S'} \to 0$ as $S' \to \infty$. Given that the European call option has payoff at maturity given by $\psi(S') := (S - S')^+$ then it is clear that the call price tends to zero as $S' \to \infty$, and hence so does the first strike derivative.

Secondly we must assume that the conditional density $p(S',\tau|S,t) \to 0$ sufficiently fast so that $S'p(S',\tau|S,t)$ tends to zero as $S' \to \infty$. Clearly this is the case for the log-normal conditional asset price probability distribution inferred by the asset price process (31). The lognormal probability distribution has exponential tails which disappear to zero as the strike increases linearly towards infinity.

With these assumptions, all the boundary terms are zero and can be ignored. Substituting (42a-42d) with $\tau = T$ into (41) gives the required result.

$\square$

The local volatility function $\sigma(S,t)$ can be fully determined from the solution to this PDE since all other terms in the equation are in terms of the given market data. However, since the problem is ill-posed there may be many possible local volatility functionals which fit the market data. To obtain an approximately arbitrage-free pricing algorithm the implied volatility (or alternatively call price data) supplied by the market must be fitted exactly. In an attempt to achieve this we apply the interpolation and extrapolation method of *cubic splines* to approximate the implied volatility surface.

Several methods exist for fitting a line or surface to discrete data, including *polynomial interpolation* methods, but for several reasons our choice favours cubic splines. Splines tend to be more stable than other methods and are easily extended to approximating in two dimensions, using *bicubic splines.* An application of cubic splines in one-dimension produces a fit which is smooth in the first derivative and continuous in the second derivative, however this is not the case for higher dimensions and only smoothness is then guaranteed. The ease with which approximations to these derivatives can be obtained is a major advantage in reducing the computational complexity of our method for determining the local volatility described in §3.3 below.

Bicubic splines extend one-dimensional cubic splines by constructing n_K row splines (one for each strike y_j) of length n_T across the volatility data. When these approximations have been constructed, the value for any strike K at maturity T can be found by constructing a column spline of length n_K down the tabulated values $v_i(K),\ i = 1, \ldots, n_K$. The order of the interpolation is chosen to ensure stability - in general data is available for many more strike prices than maturities.

The precomputed algorithm for finding the value of volatility for a general strike K and maturity T can be summarised as follows:

- Fit n_K bicubic splines of length n_T to the rows of the volatility data.
- Calculate and store the derivatives v_{TT}, v_T and the volatility v for maturity T on each row spline to obtain values for all strike prices at the given maturity T.
- Fit a spline of length n_K across the newly created values, and calculate the volatility value and derivatives with respect to the strike K for maturity T.

Since the construction of a length N cubic spline is the solution of a tridiagonal system it is an $\mathcal{O}(N)$ calculation. The complexity of the problem can be reduced by precomputing more derivative information, at the expense of additional memory requirements.

3.3 Methodology

We now outline the approach used to obtain a local volatility surface which fits the market data – in the form of implied volatilities of European call options – and uses bicubic spline interpolation to create a continuous analogue of the discrete data available in the market. Further, a consistent *local* volatility must be derived so that we can use the linear programming approach developed in §2 to price American exotic options consistent with the volatility smile.

Since we are considering the implied volatility non-constant in the underlying diffusion, assume that the quoted European call prices are Black-Scholes prices with non-constant implied volatilities as an additional parameter. Hence, define the European call prices in terms of the *implied volatility* $\nu(K,T)$ associated with the call option of strike K and maturity T. Then

$$C(K,T) := C_{BS}(S,t;K,T,\nu(K,T)),$$

where C_{BS} is the Black-Scholes European call price.

Using this formulation we can write the local volatility in terms of derivatives of the Black-Scholes implied volatility, using the following proposition [3].

[3] This result was derived following discussions with S.H. Babbs (then of HSBC Markets, now First National Bank of Chicago) and has since been independently derived in [2].

Proposition 5

Let the asset price **S** *follow the diffusion process in (31). Then the local volatility function* $\sigma(S,t)$ *consistent with the arbitrage-free European call prices is given uniquely, in the absence of dividends*[4]*, by*

$$\sigma^2(K,T) = 2\frac{\frac{\partial C}{\partial T} + r(T)K\frac{\partial C}{\partial K}}{K^2\frac{\partial^2 C}{\partial K^2}} \tag{43}$$

with $S = K$.

In terms of the implied volatility function $\nu(K,T)$ *this can be written as*

$$\sigma^2(K,T) = \frac{2\left(\sqrt{T-t}\frac{\partial \nu}{\partial T} + \frac{1}{2}\frac{\nu}{\sqrt{T-t}} + r(T)K\frac{\partial \nu}{\partial K}\right)}{K^2\left(\frac{1}{\nu}\left(\frac{\partial \nu}{\partial K}\left((T-t)-d_1\right) - \frac{1}{K\sqrt{T-t}}\right)\left(-d_1\frac{\partial \nu}{\partial K}\sqrt{T-t} - \frac{1}{K}\right) + \frac{\partial^2 \nu}{\partial K^2}\sqrt{T-t} + \frac{\sqrt{T-t}}{K}\frac{\partial \nu}{\partial K}\right)} \tag{44}$$

where

$$d_1 := \frac{\ln(S/K) + \left(r + \frac{1}{2}\nu(K,T)^2\right)(T-t)}{\nu(K,T)\sqrt{T-t}}. \tag{45}$$

Proof

Equation (43) follows immediately from (37). Following the portfolio dominance arguments in [2] it can be shown that $\sigma^2(K,T)$ is non-negative in the absence of arbitrage if the numerator of (43) is non-negative. This result follows since the denominator is a probability density which is non-negative in an arbitrage-free world.

Equation (44) follows from (43) after much calculation. Using the chain rule of differentiation we can rewrite (43) as

$$\sigma^2(K,T) = 2\frac{C_\nu \nu_T + r(T)KC_\nu \nu_K}{K^2\left(C_{K\nu}\nu_K + C_\nu \nu_{KK}\right)}. \tag{46}$$

Since C is simply the solution of the Black-Scholes equation given by (45) we can obtain expressions for the derivatives with respect to ν in (46), giving

$$C_\nu = S\mathcal{N}'(d_1)\sqrt{T-t}$$

where

$$\mathcal{N}'(x) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}x^2}$$

and

$$\begin{aligned} C_{KK} &= C_{K\nu}\nu_K + C_\nu \nu_{KK} \\ &= \frac{1}{\nu}\left(\frac{\partial \nu}{\partial K}\left((T-t)-d_1\right) - \frac{1}{K\sqrt{T-t}}\right)\left(-d_1\frac{\partial \nu}{\partial K}\sqrt{T-t} - \frac{1}{K}\right) + \frac{\partial^2 \nu}{\partial K^2}\sqrt{T-t} + \frac{\sqrt{T-t}}{K}\frac{\partial \nu}{\partial K} \end{aligned} \tag{47}$$

[4] The addition of a deterministic dividend process is a simple extension of this work.

which upon substitution into (46) gives (44) as required. $\square$

We will use this algebraic expression to obtain the local volatilities. An alternative derivation has been produced in [2] and independently in [3]. The market call prices could have been utilized directly through interpolation and the use of (43) but this method has several shortcomings. For example, the call prices for high strikes tend to zero so floating-point numerical inaccuracies occur in the derivatives of these call prices. It is also possible for the interpolated call prices to become negative for deep out-of-the-money options. Another similar approach is to directly interpolate the Black-Scholes implied volatilities and to use the Black-Scholes equation in a transitional sense to obtain the call option prices. This method also breaks down in practice since the standard normal distribution $\mathcal{N}(\cdot)$ must be numerically approximated and this is difficult to achieve efficiently to a higher accuracy than 10^{-10}. For high strikes and small maturities this method produces negative call prices and is thus an unstable algorithm. For the method followed here the implied volatilities are relatively stable and our expression for the local volatility involves relatively few numerical calculations and only one computationally expensive logarithmic calculation for each point at which it is evaluated.

First a bicubic spline interpolation is fitted to the market implied volatility data, with the calculated second derivatives with respect to the maturity stored in an array. By firstly fitting the splines across maturities for each strike we obtain approximations for the first and second derivatives with respect to the strike when the final column spline is evaluated. The first-order derivative of the volatility with respect to maturity is not a natural by-product of the interpolation but is instead approximated by a simple first order approximation. Given strike and maturity values at a mesh point the cubic spline interpolation is sufficient to supply all the values required for the calculation of the local volatility in (44).

Since the option valuation takes place on a log-transformed grid for consistency with the earlier developed methods, the strike prices need to be transformed so that the same grid can be utilized for calibration and pricing. Defining $\zeta := \ln(K)$ and $\tilde{C}(\zeta, T) := C\left(e^{\zeta}, T\right)$ the local volatility is given in terms of ζ as

$$\sigma^2(\zeta,T) = \frac{2\left(\sqrt{T-t}\,\frac{\partial\nu}{\partial T} + \frac{1}{2}\frac{\nu}{\sqrt{T-t}} + r(T)\frac{\partial\nu}{\partial\zeta}\right)}{e^{\zeta}\left(\frac{1}{\nu}\left(\frac{\partial\nu}{\partial\zeta}\left((T-t)-\tilde{d}_1\right) - \frac{1}{\sqrt{T-t}}\right)\left(-\tilde{d}_1\frac{\partial\nu}{\partial\zeta}\sqrt{T-t}-1\right) + \left(\frac{\partial^2\nu}{\partial\zeta^2} - \frac{\partial\nu}{\partial\zeta}\right)\sqrt{T-t} + e^{-\zeta}\sqrt{T-t}\,\frac{\partial\nu}{\partial\zeta}\right)} \tag{48}$$

where $\tilde{d}_1$ is simply d_1 in terms of ζ instead of K. At each node (i,m) in the grid the local volatility is calculated from the spline approximation and the array of calculated volatilities is stored to be used later in the pricing algorithm.

The pricing procedure follows from that developed for the Black-Scholes model with the modification that the volatility in the formulation is no longer considered to be constant. This changes the constraint matrix in the order complementarity problem, and thus in the linear programming formulation; the matrix now has non-constant diagonals, but is still tridiagonal in nature. However, for American and exotic options this is the ideal problem to be solved by the non-constant tridiagonal simplex method described in §2.2. Since this method calculates the UL basis decomposition at each time-step, updating the decomposition only when new basic variables enter the basis, it is a very efficient solution procedure for this problem.

3.4 Dealing with Market Arbitrage

As noted above, the main problem with fitting the volatility smile to the market data is that the result may not be arbitrage free. Several studies (e.g. [1, 11]) have aimed to detail these instabilities in the data and why they occur. If a model produces arbitrage opportunities we must use some regularization or filtering procedure applied to the original market data to remove them. Some practitioners suggest setting any negative probability densities to zero (implying an infinite local volatility), or restricting any local volatilities to lie within a range $(\sigma_{\min}, \sigma_{\max})$ where the bounds are supplied somewhat arbitrarily. In our procedure we try to filter the data using methods derived from the underlying theory.

When a negative denominator occurs in (44) we necessarily have an arbitrage opportunity since this implies a negative value of the transitional probability density. We correct this value using put-call parity [33] to consider the prices of European put options implied by the call option market data. Inherent in the finite difference solution is the approximation of an integral by a discrete sum, namely

$$\int_0^{\infty} \psi(K) p(K, T|S, t) dK \approx \sum_{i=0}^{N} \psi(K_i) p(K_i, T|S, t) \Delta K, \tag{49}$$

where $p(\cdot) := p(\cdot, T|S, t)$ now denotes the discrete form of the conditional density. If the market data implies a negative value of $p(K_j)$ for some $j \in [0, N]$ we consider a European put option of strike K_{j+1} with value $P(S, t; K_{j+1}, T)$. Then

$$\begin{aligned} P(S, t; K_{j+1}, T) &\approx e^{-r(T-t)} \sum_{i=0}^{j} (K_{j+1} - K_i) p(K_i) \Delta K \\ &= e^{-r(T-t)} \sum_{i=0}^{j-1} (K_{j+1} - K_i) p(K_i) \Delta K + e^{-r(T-t)} (K_{j+1} - K_j) p(K_j) \Delta K \end{aligned} \tag{50}$$

which given all European put option prices are known (from the assumption that all European call option prices are known together with put-call parity) allows us to find a value of the discrete probability $p(K_j)$ consistent with the market data. We must of course ensure that this probability is non-negative, otherwise its value is set to zero.

The other possible inconsistency in the data is when the numerator of (44) is negative. We rectify this by setting the local volatility at this point to zero.

4 Numerical Results

In this section we give empirical results for the procedure for pricing options consistent with the observed market volatility smile outlined in §3.3. To make comparisons with published articles we use two sets of real market data previously utilized as our underlying implied volatility surfaces and price European, American and exotic options with respect to this data. All solution times quoted are for results calculated on an IBM RS6000/590 workstation with 1 GB RAM running under AIX 4.3, although only a small proportion of this memory is utilized. Results are quoted in [20, 44] for solution on a Pentium II 400Mhz machine running under Windows NT 4, which gives significant speed-ups for most levels of domain discretization.

4.1 Volatility Data

Several approaches in the literature 'manufacture' a volatility surface as a test data set, but to test the numerical procedure described in §3.3 we use two different market data sets as noted above. The first contains European call option prices quoted on the S&P 500 index as used in [2, 13] and the second consists of FTSE 100 volatility values implied from European call option data as described in [26].

The corresponding two sets of implied volatility data are given in Tables 1 and 2 and graphically in Figures 2 and 3. The data in Table 1 refers to the Black-Scholes implied volatilities of European call options on the S&P 500 index as observed in October 1995. At the time the data was gathered the S&P 500 index stood at a level of 590, and European call options were available for a range of strikes from 85-140% of the current price, an index range of 501.50 to 826.00. Call options were also available for maturities ranging from 0.175-5 years, but as in [2] this study uses only short-maturity data with $T < 2$ years. It should be noted that the data given is not all quoted in the market as the volatilities have been extrapolated in areas of the table for which data was not available. The authors of [2] have extrapolated a severe upward sloping section of the data for short maturities and high strikes to put additional pressure on the pricing algorithm. For benchmark results we use only options with maturity less than the mid-range of the data, i.e. with $T \leq 1$ year, since this data is the most reliable. Figure 2 illustrates the implied volatility test surface given by the data in Table 1 and shows the dependence of the implied volatility on the maturity and strike price, highlighting smile and temporal volatility effects.

Table 2 corresponds to FTSE 100 volatility values implied from European call option data for 31 March 1995. The initial index level was 3129.5 and data was quoted for 8 strike prices and 5 maturities, but prices were quoted for different strikes for the last two maturities than for the earlier three maturities. However the data was simply interpolated to fill the gaps as described in §3.2. We also assume that a constant rate of interest $r = 10\%$ applies throughout the period from time 0 until maturity 0.737 years.

	Maturity									
Strike	0.175	0.425	0.695	0.940	1.000	1.500	2.000	3.000	4.000	5.000
501.50	0.190	0.177	0.172	0.171	0.171	0.169	0.169	0.168	0.168	0.168
531.00	0.168	0.155	0.157	0.159	0.159	0.160	0.161	0.161	0.162	0.164
560.50	0.133	0.138	0.144	0.149	0.150	0.151	0.153	0.155	0.157	0.159
590.00	0.113	0.125	0.133	0.137	0.138	0.142	0.145	0.149	0.152	0.154
619.50	0.102	0.109	0.118	0.127	0.128	0.133	0.137	0.143	0.148	0.151
649.00	0.097	0.103	0.104	0.113	0.115	0.124	0.130	0.137	0.143	0.148
678.50	0.120	0.100	0.100	0.106	0.107	0.119	0.126	0.133	0.139	0.144
708.00	0.142	0.114	0.101	0.103	0.103	0.113	0.119	0.128	0.135	0.140
767.00	0.169	0.130	0.108	0.100	0.099	0.107	0.115	0.124	0.130	0.136
826.00	0.200	0.150	0.124	0.110	0.108	0.102	0.111	0.123	0.128	0.132

Table 1: Implied volatilities of S&P 500 equity index European call options for October 1995 - Initial index 590

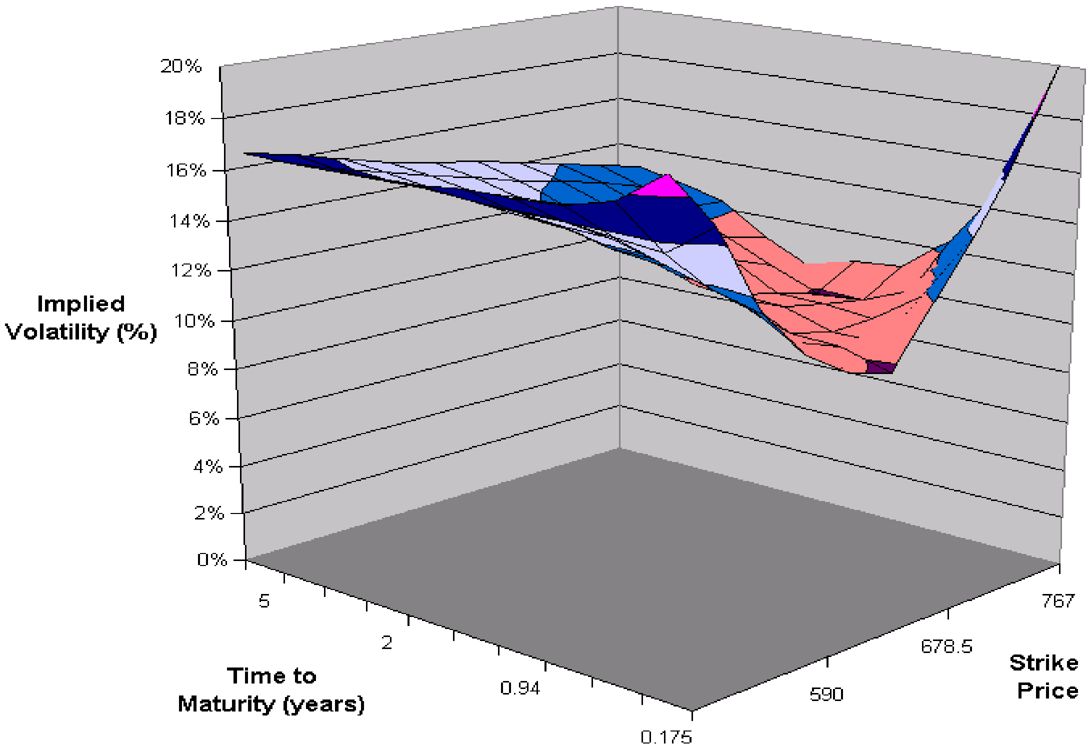

Figure 2: Implied volatility surface for European call options on the S&P 500 index

Strike Price	April	May	June	Sept	Dec
2975	17.36%	15.92%	16.52%		
3025	16.16%	15.57%	16.04%	16.25%	16.23%
3075	14.67%	15.22%	15.72%		
3125	14.45%	14.79%	15.41%	15.69%	15.89%
3175	14.30%	14.46%	15.05%		
3225	13.64%	14.16%	14.61%	14.97%	15.46%
3275	13.94%	13.98%	14.25%		
3325	13.21%	13.75%	13.83%	14.39%	14.92%
Maturity in Years	0.057534	0.134247	0.210959	0.460274	0.736986

Table 2: Implied volatilities for the FTSE 100 equity index European call options for March 31 1995 - Initial index 3129.5

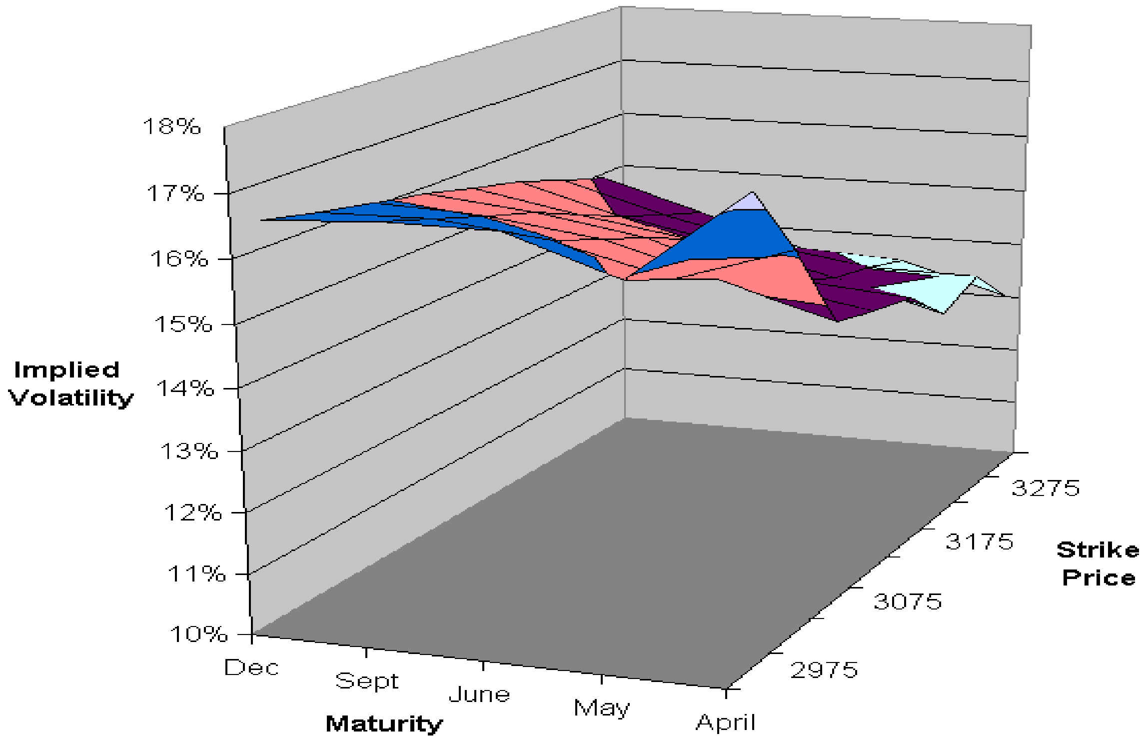

Figure 3: Implied volatility surface for European call options on the FTSE 100 index

4.2 Computational Procedure

As described in §3.3 the pricing procedure occurs in two stand-alone stages. The first - calibrating the local volatility - utilises the estimates of the implied volatility derivatives to calculate the surface $\sigma(S,t)$ for use in pricing. In the following results this calibration takes place on the same mesh which is used in the pricing procedure. However there are significant computational savings to be made by only calculating the local volatilities on a coarser mesh and performing some type of interpolation between the calibrated values; this is not investigated further in this paper.

The initial calibration of the surface is computationally expensive. To calculate a volatility surface for a discretization of 800 time-steps and 200 space-steps takes approximately 20 seconds, although the time required for subsequent valuations on the same grid is a fraction of this if the surface is stored for future use. We only quote solution times for the option valuation, assuming that the initial creation of the local volatility surface is in memory.

After calibrating the volatility surface the options are valued using the UL update algorithm of §2 for options with American exercise and for European options using a simple linear equation solver. For the former, at each time-step of the UL update the basis decomposition was calculated and the update used only when new variables entered the basis.

4.3 Underlying Model Accuracy

We will attempt to evaluate the accuracy of the calibration of the volatility surface using the pricing procedures previously outlined. But before doing this it is necessary to understand the effect of the numerical approximations made in the finite-difference scheme. As discussed in [20] Crank-Nicolson finite-difference approximations are accurate to $\mathcal{O}((\Delta t)^2 + (\Delta x)^2)$ so that this underlying numerical error – *not* due to the volatility calibration – must be accounted for in any comparisons between the given market data and computed option prices. To enable this comparison to be made, the values of European call options on the S&P 500 index (Table 1) of less than 1 year maturity were calculated using the Crank-Nicolson scheme (described in [20]) with a reasonable discretization level and σ was set to the at-the-money Black-Scholes implied volatility σ_{atm} for each strike K. The *tridag* solution algorithm described in [43] replaced the tridiagonal LP solver, since for European options the discretized PDE is simply a linear vector equation. These results are displayed in Table 3 and show that for quite a low discretization level, which is convergent to 2 decimal places, the error of the computational pricing scheme is less than 4 basis points (at the money) when compared with the market values. These benchmark market values were calculated from the Black-Scholes call pricing equation, and thus have themselves a small numerical error due to approximating the normal distribution function. Since the actual market values would be contained within a bid-ask spread, this accuracy is acceptable and is our benchmark for the accuracy of the volatility fit.

4.4 European Option Results

With an understanding of the underlying numerical pricing accuracy of the finite-difference approach, we now use our pricing algorithm to recover the values of all the options given in the market in order to assess the goodness-of-fit of the calibrated surface.

Since the S&P 500 volatility data is only used for maturities less than 2 years the results in Table 4 concentrate on shorter maturity European call options. As discussed in [2] the validity of option prices cannot be guaranteed for longer maturities since the volume of trade in these markets is much less than for shorter maturities. Table 4 contains European call option prices for the strikes and maturities given by the market. The reference values are calculated through the use of the Black-Scholes pricing equation [6] with implied volatility values given by Table 2. This requires an approximation for the cumulative Normal distribution which is accurate to 10 decimal places [33]. The results in this table are for the same parameters as in Table 3 using the same procedure, so that an accuracy comparison can be made. The errors can be seen to be comparable to - more accurate in some cases than - the baseline numerical accuracy described in Table 3. The conclusion must be that fitting the volatility smile does not induce any significant errors above the baseline accuracy into the option values for European call options. Since the original European call options can be seen to be accurately priced, the conclusion must be that the volatility surface that is fit to the data is consistent with the local volatilities implied by the market through the prices it quotes. Solution time for a European option at this discretization is approximately 0.4 seconds. The actual calculated local volatility surface can be seen in Figure 4. A comparison with the implied volatility surface shown in Figure 2 shows significant differences.

The local volatility surface (Figure 4) is represented on a truncated strike domain. At short maturities a spike of local volatility occurs at a strike value of approximately 650 which

Maturity (years)	Strike	Implied volatility	Market value	PDE Value	Error ×100
	501.5	0.177	110.296	110.296	0.012
	531	0.155	82.857	82.856	0.054
	560.5	0.138	56.708	56.694	1.432
	590	**0.125**	**33.547**	**33.510**	**3.720**
0.425	619.5	0.109	14.979	14.962	1.708
	649	0.103	5.006	5.037	3.135
	678.5	0.1	1.205	1.224	1.890
	708	0.114	0.545	0.562	1.707
	767	0.13	0.083	0.088	0.487
	826	0.15	0.025	0.027	0.184
	501.5	0.172	123.845	123.842	0.306
	531	0.157	97.630	97.627	0.256
	560.5	0.144	72.615	72.602	1.345
	590	0.133	49.804	49.776	2.818
0.695	619.5	0.118	29.460	29.441	1.891
	649	0.104	13.863	13.878	1.495
	678.5	0.1	5.672	5.682	0.965
	708	0.101	2.196	2.224	2.830
	767	0.108	0.336	0.349	1.265
	826	0.124	0.108	0.113	0.530
	501.5	0.171	138.610	138.605	0.437
	531	0.159	113.447	113.442	0.453
	560.5	0.15	89.551	89.539	1.248
	590	0.138	66.743	66.720	2.327
1	619.5	0.128	46.397	46.380	1.786
	649	0.115	28.393	28.394	0.153
	678.5	0.107	15.469	15.462	0.696
	708	0.103	7.616	7.640	2.365
	767	0.099	1.338	1.363	2.532
	826	0.108	0.361	0.373	1.219

Table 3: Pricing Error: At the money European call options. $T \leq 1$ years, $M = 800$, $I = 200$.

M	I	K	T	BS σ	Value	BS Value	Error($\times$100)
800	200	501.5	0.175	19.00%	97.37	97.37	0.81
800	200	531	0.175	16.80%	68.82	68.84	2.33
800	200	560.5	0.175	13.30%	40.89	40.87	1.42
800	200	590	0.175	11.30%	16.90	16.89	0.85
800	200	619.5	0.175	10.20%	3.48	3.45	2.57
800	200	649	0.175	9.70%	0.28	0.26	1.90
800	200	678.5	0.175	12.00%	0.04	0.08	4.01
800	200	708	0.175	14.20%	0.00	0.03	2.66
800	200	767	0.175	16.90%	0.00	0.00	0.29
800	200	826	0.175	20.00%	0.00	0.00	0.09
800	200	501.5	0.425	17.70%	110.26	110.30	3.14
800	200	531	0.425	15.50%	82.83	82.86	2.79
800	200	560.5	0.425	13.80%	56.68	56.71	3.14
800	200	590	0.425	12.50%	33.51	33.55	3.91
800	200	619.5	0.425	10.90%	14.98	14.98	0.34
800	200	649	0.425	10.30%	5.01	5.01	0.45
800	200	678.5	0.425	10.00%	1.18	1.20	2.30
800	200	708	0.425	11.40%	0.32	0.55	2.56
800	200	767	0.425	13.00%	0.03	0.08	5.00
800	200	826	0.425	15.00%	0.01	0.02	1.52
800	200	501.5	0.695	17.20%	123.80	123.85	4.74
800	200	531	0.695	15.70%	97.58	97.63	4.76
800	200	560.5	0.695	14.40%	72.57	72.62	4.99
800	200	590	0.695	13.30%	49.75	49.80	5.53
800	200	619.5	0.695	11.80%	29.42	29.46	3.86
800	200	649	0.695	10.40%	13.91	13.86	4.43
800	200	678.5	0.695	10.00%	5.67	5.67	0.61
800	200	708	0.695	10.10%	2.19	2.20	0.52
800	200	767	0.695	10.80%	0.34	0.34	0.34
800	200	826	0.695	12.40%	0.11	0.11	0.12
800	200	501.5	1	17.10%	138.54	138.61	6.75
800	200	531	1	15.90%	113.38	113.45	6.43
800	200	560.5	1	15.00%	**89.46**	**89.55**	**9.42**
800	200	590	1	13.80%	66.66	66.74	7.91
800	200	619.5	1	12.80%	46.32	46.40	7.45
800	200	649	1	11.50%	28.39	28.39	0.38
800	200	678.5	1	10.70%	15.44	15.47	2.72
800	200	708	1	10.30%	7.62	7.62	0.80
800	200	767	1	9.90%	1.36	1.34	1.81
800	200	826	1	10.80%	0.37	0.36	0.49

Table 4: European call option prices fitting the S&P 500 volatility - short maturities. Initial spot $S = 590$. Value is the calculated value. BS value is the Black-Scholes price using the given implied volatility.

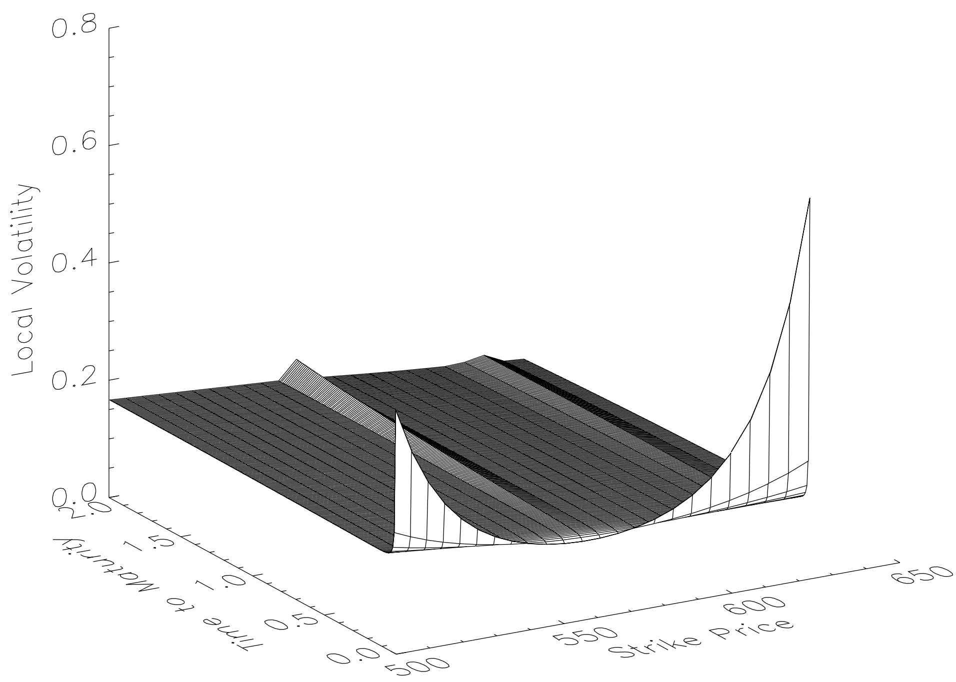

Figure 4: Local volatilities $\sigma(S,t)$ for S&P 500 index options on a truncated strike domain. Parameters $M = 200$, $I = 200$ for the full strike range.

distorts any graphical representation of the local volatility surface, but does not cause any instability in the calculated option values. For ease of representation, the local volatility is shown for all strikes less than the level at which the spike occurs.

Results for the FTSE 100 data show similar patterns. Table 5 contains specific pricing errors of European call options for different levels of the domain discretization and shows the convergence properties of the algorithm. We compare the calculated value with the value given by the Black-Scholes call pricing equation to define the quoted error. As can be seen the error is behaving properly, with a reduction in the error as both time and space discretizations increase. This particular option had value 114.194 so an error of ≈ 0.115 is equivalent to a 0.1% error in the option value. It can be seen that the number of time-steps is quite critical and this may be due to the discontinuity in the interpolation of the term structure of the volatility. The derivative of volatility with respect to maturity is approximated by a first-order difference since this term is not a natural by-product of bicubic spline interpolation. However, a discretization level of $M = 1600$ and $I \approx 300$ gave an error of less than 0.1%. If we require accuracy to within 0.1% of strike then all levels of discretization would satisfy

		Number of Time-Steps					
		100	200	400	800	1600	3200
Space	100	-0.867	-0.370	-0.114	0.020	0.097	0.141
Steps	200	-0.998	-0.508	-0.262	-0.138	-0.074	-0.040
	400	-0.984	-0.494	-0.249	-0.127	-0.065	-0.034

Table 5: Errors in the calculated values of a European option on the FTSE 100 index with changes in the discretization level. Strike $K = 3125$ $T = 0.211$ $r = 0.1$ $S = 3109.5$.

Strike	April			May			June		
Price	Value	Actual	Error	Value	Actual	Error	Value	Actual	Error
2975	158.84	158.87	3.50	189.08	189.09	1.00	221.23	221.25	2.30
3025	114.92	114.88	4.40	148.82	148.80	1.80	181.84	181.82	1.70
3075	74.54	74.52	1.63	112.51	112.52	0.70	146.10	146.11	1.20
3125	44.43	44.46	2.89	80.69	80.72	3.32	114.15	114.19	4.00
3175	23.50	23.51	1.09	54.92	54.91	0.54	86.11	86.13	1.93
3225	9.77	9.79	2.34	**35.13**	**35.08**	**4.88**	62.02	61.98	3.99
3275	4.15	4.17	2.10	21.24	21.21	2.62	42.85	42.84	0.94
3325	1.09	1.06	2.79	11.85	11.82	3.48	27.83	27.80	3.36

Table 6: Pricing errors $\times 100$ of FTSE 100 European call options. Actual value given by Black-Scholes call pricing equation. $M = 3200$, $I = 200$.

this criterion. Table 6 shows the pricing error for original quoted FTSE 100 options, using a discretization level of 3200 time-steps and 200 space-steps. All pricing errors are less than 5 basis points at this level of discretization.

The calculated local volatility surface is shown in Figure 5. The axes are truncated in a similar manner to the S&P local volatility surface, as a spike of volatility is also in evidence for the FTSE data at a strike level of approximately 3250. The figure also shows that for short maturities the local volatility is behaving like the reciprocal of the conditional probability density, smoothing out somewhat for higher maturities.

4.5 American Option Results

The next stage of our analysis is to extend the work to the pricing of American options. This introduces an added dimension to the problem, since there is a very real possibility that the optimal exercise boundary (see §2.1) will be moved by fluctuations in the local volatility surface. The actual solution is a modification of the valuation of European options in the previous section, with the updating tridiagonal LP solver introduced in place of the *tridag* linear equation solver. Boundary conditions used in this section are as described in [20].

The first results on the S&P 500 in Table 7 show American put valuations for the LP tridiagonal solver and also for the PSOR algorithm described in [14]. As a benchmark we use the original LP valuation on a very fine solution mesh using the at-the-money BS implied volatility. For the American option we can immediately see that the values are *lower* than

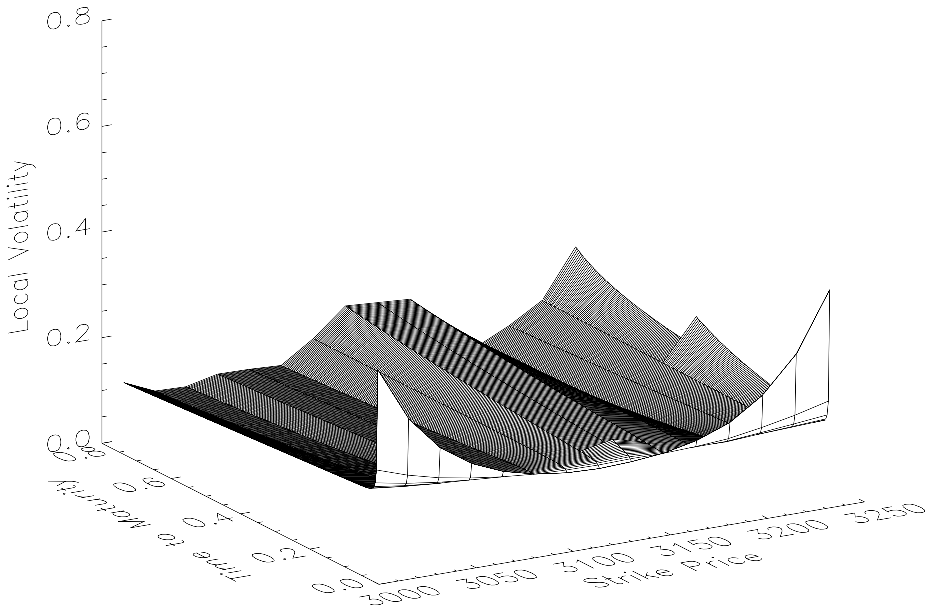

Figure 5: Local volatilities $\sigma(S, t)$ for FTSE 100 index options on a truncated strike domain. Parameters $M = 200$, $I = 200$ for the full strike range.

the constant volatility values to within the numerical tolerance allowed. The PSOR and LP solutions both have comparable values on a relatively small grid, however the LP solution time is 0.4 seconds whilst the PSOR solution time is close to 9 seconds. Since both solution algorithms converge to within the same accuracy, the discrepancy between the smile-fitting and at-the-money implied values must be due to the volatility surface.

It was noted in §2 that the convex shape of the optimal exercise boundary for the American put problem could be used to increase the efficiency of the tridiagonal solver. Figure 6 highlights the reasons why the option value is so different when the volatility smile is fitted by illustrating the difference in the shape of the optimal exercise boundary for the option. When we take account of the local volatility, the exercise boundary is no longer a smooth function of the asset price as might be expected, but is shifted by changes in the volatility. Whilst this is no problem for the accuracy of the pricing algorithm, it does radically affect the realised option price.

Table 8 contains comparison results for the FTSE 100 index. For this data the smile-fitting values are significantly *higher* than the LP(σ_{LP}) constant volatility computed values, reflecting the opposite slope of the short-dated option 'smirk' in this case to that of the S&P 500 options.

K	σ	LP	PSOR	LP(σ_{atm})
501.5	17.100%	2.695	2.696	2.704
531	15.900%	4.630	4.633	4.776
560.5	15.000%	8.378	8.393	8.946
590	13.800%	15.244	15.301	16.137
619.5	12.800%	30.230	30.352	31.145
649	11.500%	59.000	59.000	59.000
678.5	10.700%	88.500	88.500	88.500
708	10.300%	118.000	118.000	118.000
767	9.900%	177.000	177.000	177.000
826	10.800%	236.000	236.000	236.000

Table 7: American put valuation results fitting the S&P 500 volatility smile. $T = 1$ year, $M = 800$, $I = 200$. PSOR tolerance$=1\times10^{-8}$, $\omega = 1.2$. LP(σ_{imp}) corresponds to the 1-factor tridiagonal solution using at-the-money BS implied volatility with $M = 10000$, $I = 10000$.

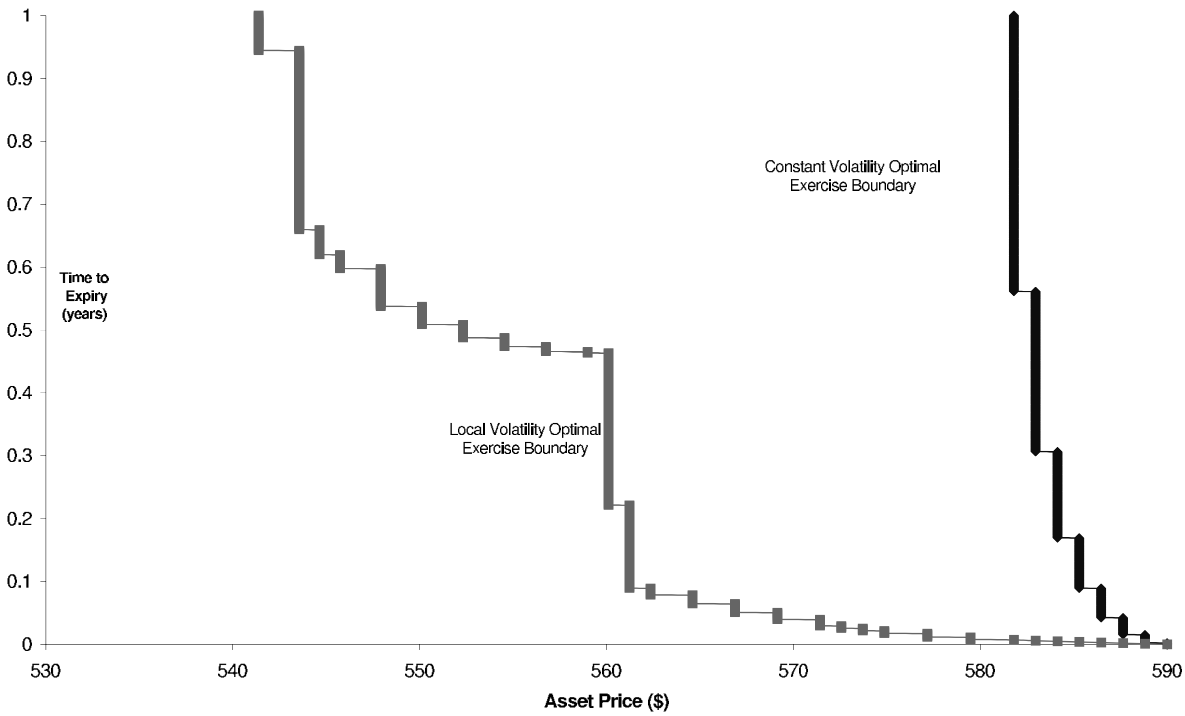

Figure 6: American put optimal exercise boundary for the smile fitting LP solution. Comparison with the 1-factor LP method using ATM BS implied volatility. Parameters: $M = 1000, I = 1000$ $S_0 = K = 590$, $\sigma_{BS} = 13.8\%$

Strike	April		May		June	
Price	Smile Fit	LP(σ_{atm})	Smile Fit	LP(σ_{atm})	Smile Fit	LP(σ_{atm})
2975	9.31	7.12	18.54	15.01	29.86	24.62
3025	15.72	12.77	28.66	24.05	40.69	34.13
3075	26.03	21.86	42.92	37.11	55.40	47.43
3125	46.40	41.77	62.28	54.65	74.50	64.45
3175	76.53	70.38	88.59	78.01	98.70	85.33
3225	116.51	106.80	122.79	107.71	129.19	110.15
3275	166.50	150.90	165.84	143.13	167.70	139.95
3325	215.50	197.59	215.50	183.13	215.50	173.86

Table 8: American put option valuation results fitting the FTSE 100 volatility smile. LP(σ_{atm}) is the value calculated from the tridiagonal solver with discretization $M = 10000$, $I = 10000$.

4.6 Exotic Option Results

We now price Asian fixed-strike options by fitting the smile. A discretely-sampled *Asian fixed-strike* option's value is dependent on a pre-specified strike value K and the current value of a discretely-sampled arithmetic average of the asset price, where the average is calculated on certain sample dates.

All results correspond to at-the-money options of 1 year maturity with the risk-free rate assumed constant at 10%. In all the tables the column *implied value* refers to the value found using the LP approach with constant volatility set to the Black-Scholes at-the-money implied volatility for the options in question. The column *smile value* illustrates the results obtained by fitting the volatility smile and term structure.

Table 9 contains option values for the European Asian put option with fixed-strike equal to the initial asset price. As can be seen the option price fitting the volatility smile is significantly greater than the constant volatility price. This is the pattern seen throughout all S&P 500 put option results for all sampling levels, with a similar effect being seen in the price of call options. In general, all these options are priced higher by fitting the smile, despite the fact that the same volatility structure was used to accurately fit the European call option prices in §4.4.

The corresponding American option values are shown in Table 10. Unlike the vanilla American put, the modelling of the volatility *adds*, rather than subtracts, a premium from the option value for the Asian option. The solution time for a discretization of 200×200^2 is approximately 18 seconds.

Conversely, the American fixed-strike Asian option results on the FTSE 100 index (Table 11) show a mixed effect. For a small number of samples of the average the smile-fitting value is less than the Black-Scholes implied value, but the converse relationship holds for higher sampling rates. This effect is likely due to the time-to-maturity variability of implied volatility at strike 2975 depicted in Figure 3.

Samples	M	I	J	Implied Value	Smile Value
	200	100	100	5.12	6.32
2	200	200	200	5.16	6.35
(6 monthly)	400	200	200	5.16	6.38
	260	100	100	6.81	8.60
12	260	200	200	6.80	8.58
(monthly)	520	200	200	6.80	8.58

Table 9: Discretely sampled European fixed-strike Asian put option results fitting the S&P 500 smile. Parameters: $K = 590$, $S_0 = 590$ and $T = 1$ year.

Samples	M	I	J	Implied Value	Smile Value
	200	100	100	5.38	6.64
2	200	200	200	5.42	6.67
(6 monthly)	400	200	200	5.42	6.68
	260	100	100	9.71	12.03
12	260	200	200	10.20	12.57
(monthly)	520	200	200	10.20	12.67

Table 10: Discretely sampled American fixed-strike Asian put option results fitting the S&P 500 smile. Parameters: $K = 590$, $S_0 = 590$ and $T = 1$ year

SAMPLES	M	I	J	Implied Value	Smile Value
	400	200	200	0.360	0.295
2	800	200	200	0.360	0.296
(6 monthly)	800	400	400	0.353	0.278
	270	200	200	3.134	3.710
12	540	200	200	3.135	3.759
(monthly)	540	400	400	3.127	3.737

Table 11: Discretely sampled American fixed-strike Asian put option results fitting the FTSE 100 smile. Parameters: $K = 2975$, $S_0 = 3137$ and $T = 0.211$ years.

5 Conclusions and Future Directions

In this paper we have applied a fast accurate linear programming valuation algorithm to pricing exotic American options fitting the volatility smile implied by the market prices of vanilla European call options. We have demonstrated first that the basic Crank-Nicolson finite difference methods have low discretization error and that the quoted vanilla options are accurately priced by the fitted local volatility surface. Subsequently we have seen that due to local volatility effects on the computed optimal exercise boundary, prices of American options fitted to the smile differ significantly from those with constant volatilities. Finally, we have seen similar effects for exotic options – as represented by discretely sampled fixed-strike Asian options.

Current research extends the testing of these methods to lookbacks and barriers, including both digitals and knock-in and knock-out features for Asians and lookbacks. An interesting area of related research involves the Kalman filtering of local volatility surfaces – as for example computed in this paper – from one market epoch (day) to the next in order to achieve better long-run hedging. Another line of our current research with PDE-based valuation methods concerns wavelet basis techniques for high-dimensional Bermudan and American fixed income derivatives.

References

[1] Amin, K. I. and V. K. Ng (1997). Inferring Future Volatility from the Information in Implied Volatility in Eurodollar Options: A New Approach. *Review of Financial Studies* **10** (2) 333–367.

[2] Andersen, L. B. G. and R. Brotherton-Ratcliffe (1997). The Equity Option Volatility Smile: An Implicit Finite-Difference Approach. *Journal of Computational Finance* **1** (2) 5–37.

[3] Andreasen, J. (1996). Implied Modelling: Stable Implementation, Hedging and Duality. Working Paper, University of Aarhus.

[4] Ball, C. A. and A. Roma (1994). Stochastic Volatility Option Pricing. *Journal of Financial and Quantitative Analysis* **29** (4) 589–607.

[5] Barraquand, J. and T. Pudet (1996). Pricing of American Path-Dependent Contingent Claims. *Mathematical Finance* **6** (1) 17–51.

[6] Black, F. and M. Scholes (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* **81** 637–659.

[7] Bodurtha, J., Jr (1997). A Linearization-Based Solution to the Ill-Posed Local Volatility Estimation Problem. Working Paper, Georgetown University.

[8] Bodurtha, J., Jr and M. Jermakyan (August 1996). Non-Parametric Estimation of an Implied Volatility Surface. IAFE Computational Finance Proceedings, Stanford University, August 23.

[9] Borwein, J. M. and M. A. H. Dempster (1989). The Linear Order Complementarity Problem. *Mathematics of Operations Research* **14** (3) 534–558.

[10] Breeden, D. and R. Litzenberger (1978). Prices of State-Contingent Claims Implicit in Option Prices. *Journal of Business* (51) 621–651.

[11] Canina, L. and S. Figlewski (1993). The Informational Content of Implied Volatility. *Review of Financial Studies* **6** (3).

[12] Clarke, N. and K. Parrott (1998). Multigrid for American Option Pricing with Stochastic Volatility. Working Paper, Department of Computer Science, University of Greenwich.

[13] Coleman, T. F., Y. Li and A. Verma (1998). Reconstructing the Unknown Local Volatility Function. Working Paper, Computer Science Department, Cornell University. Presented at the First Annual Conference Computational and Quantitative Finance '98, New York.

[14] Cryer, C. W. (1971). The Solution of a Quadratic Programme Using Systematic Overrelaxation. *SIAM Journal of Control and Optimization* **9** 385–392.

[15] Cryer, C. W. and M. A. H. Dempster (January 1980). Equivalence of Linear Complementarity Problems and Linear Programs in Vector Lattice Hilbert Spaces. *SIAM Journal of Control and Optimization*

18 (1) 76–90.
[16] Dempster, M. and S. R. Pliska, eds. (1997). *Mathematics of Derivative Securities.* Cambridge University Press.
[17] Dempster, M. A. H. and J. P. Hutton (1997). Fast Numerical Valuation of American, Exotic and Complex Options. *Applied Mathematical Finance* **4** (1) 1–20.
[18] Dempster, M. A. H. and J. P. Hutton (1997). Numerical Valuation of Cross-Currency Swaps and Swaptions. In Dempster and Pliska [16], 473–503.
[19] Dempster, M. A. H. and J. P. Hutton (1999). Pricing American Stock Options by Linear Programming. *Mathematical Finance* **9** (3) 229–254.
[20] Dempster, M. A. H., J. P. Hutton and D. G. Richards (1998). LP Valuation of Exotic American Options Exploiting Structure. *Journal of Computational Finance* **2** (1) 61–84.
[21] Derman, E. and I. Kani (1994). Riding on a Smile. *Risk* **7** (2) 32–39.
[22] Derman, E. and I. Kani (1994). The Volatility Smile and Its Implied Tree. Goldman Sachs Quantitative Strategies Research Notes.
[23] Derman, E. and I. Kani (1998). Stochastic Implied Trees: Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility. *International Journal of Theoretical and Applied Finance* **1** (1) 61–110.
[24] Derman, E., I. Kani and N. Chriss (1996). Implied Trinomial Trees of the Volatility Smile. *Journal of Derivatives* **3** 7–22.
[25] Dewynne, J. and P. Wilmott (March 1993). Partial to the Exotic. *Risk* **6** (3) 38–46.
[26] Duan, J.-C. (1995). The Garch Option Pricing Model. *Journal of Finance* **5** 13–32.
[27] Duffie, D. (1992). *Dynamic Asset Pricing Theory.* Princeton University Press, Princeton, New Jersey.
[28] Dupire, B. (1994). Pricing with a Smile. *Risk* **7** (1) 18–20.
[29] Dupire, B. (1997). Pricing and Hedging with Smiles. In Dempster and Pliska [16].
[30] Harrison, J. M. and D. Kreps (1979). Martingales and Arbitrage in Multi-Period Security Markets. *Journal of Economic Theory* **20** 381–408.
[31] Harrison, J. M. and S. Pliska (1981). Martingales and Stochastic Integrals in the Theory of Continuous Trading. *Stochastic Processes and their Applications* **11** 215–260.
[32] Hull, J. and A. White (1987). The Pricing of Options on Assets with Stochastic Volatilities. *Journal of Financial and Quantitative Analysis* **3** 281–300.
[33] Hull, J. C. (1993). *Options, Futures and Other Derivative Securities.* Prentice-Hall International Inc, 2nd edition.
[34] Hutton, J. P. (July 1995). Fast Pricing of Derivative Securities. Ph.D. Thesis, Department of Mathematics, University of Essex.
[35] Jackwerth, J. C. and M. Rubinstein (1996). Recovering Probability Distributions from Option Prices. *Journal of Finance* **51** (5) 62–74.
[36] Jaillet, P., D. Lamberton and B. Lapeyre (1990). Variational Inequalities and the Pricing of American Options. *Acta Applicandae Mathematicae* **21** 263–289.
[37] Jamshidian, F. (1991). Bond and Option Evaluation in the Gaussian Interest Rate Model. *Research in Finance* **9** 131–170.
[38] Kirsch, A. (1996). *An Introduction to the Mathematical Theory of Inverse Problems.* Springer-Verlang New York.
[39] Lagnado, R. and S. Osher (1997). A Technique for Calibrating Derivative Security Pricing Models: Numerical Solution of an Inverse Problem. *Journal of Computational Finance* **1** (1) 13–25.
[40] Levin, A. (1998). Recovering Implied Volatility and Distributions from American Futures Option Prices Using the Regularization Method. Presented at the First Annual Conference Computational and Quantitative Finance '98, New York.
[41] Luenberger, D. G. (1984). *Linear and Nonlinear Programming.* Addison-Wesley, Reading, Mass., 2nd edition.
[42] Oskendal, B. (1995). *Stochastic Differential Equations - An Introduction with Applications.* Springer-Verlang, Berlin, 4th edition.
[43] Press, W. H., S. A. Teukolsky, W. T. Vetterling and B. P. Flannery (1992). *Numerical Recipes in C.* Cambridge University Press, 2nd edition.
[44] Richards, D. G. (1999). Pricing American Exotic Options. Ph.D. Thesis, Judge Institute of Management Studies, University of Cambridge.
[45] Rubinstein, M. (1994). Implied Binomial Trees. *Journal of Finance* **49** (3) 771–818.
[46] Rubinstein, M. (1995). As Simple as 1-2-3. *Risk* **8** (1) 44–47.
[47] Shaw, W. (1998). *Modelling Financial Derivatives with Mathematica 3.* Cambridge University Press.

Chapter 17.

[48] Shimko, D. (1993). Bounds of Probability. *Risk* (6) 33–37.

[49] Skiadopolous, G., S. Hodges and L. Clewlow (1998). The Dynamics of Implied Volatility Surfaces. Working Paper, FORC, University of Warwick.

[50] Skiadopolous, G., S. Hodges and L. Clewlow (1998). The Dynamics of Smiles. Working Paper, FORC, University of Warwick.

[51] Stephenson, G. and P. H. Radmore (1990). *Advanced Mathematical Methods for Engineering and Science Students.* Cambridge University Press.

[52] Wilmott, P., S. Howison and J. Dewynne (1993). *Option Pricing: Mathematical Models and Computation.* Oxford Financial Press.