Mathematical Finance, Vol. 10, No. 2 (April 2000), 157–177

PRICING AMERICAN OPTIONS FITTING THE SMILE

M. A. H. DEMPSTER and D. G. RICHARDS

University of Cambridge

This paper is a compendium of results—theoretical and computational—from a series of recent papers developing a new American option valuation technique based on linear programming (LP). Some further computational results are included for completeness. A proof of the basic analytical theorem is given, as is the analysis needed to solve the inverse problem of determining local (one-factor) volatility from market data. The ideas behind a fast accurate revised simplex method, whose performance is linear in time and space discretizations, are described and the practicalities of fitting the volatility smile are discussed. Numerical results are presented which show the LP valuation technique to be extremely fast—lattice speed with PDE accuracy. American options valued in the paper range from vanilla, through exotic with constant volatility, to exotic options fitting the volatility smile.

KEY WORDS: American options, exotic options, linear programming, volatility smile, inverse problem

1. INTRODUCTION

This article surveys the linear programming (LP) approach to fast valuation of American options 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 partial differential equation, a novel basis factorization, and the nature of the optimal exercise boundary, to create a pricing algorithm that is essentially *linear* in the number of discretization steps in space or time with the other held fixed. The method is applicable to a variable coefficient version of the Black–Scholes equation which allows the treatment and rapid solution of path-dependent exotic options taking account of local volatility. One emphasis of the paper is a complete overview of the theory underlying the LP approach to these difficult problems. Another is a representative survey of the numerical results obtained to date.

The paper is structured as follows. In Section 2 we review the formulation of the American put option valuation problem presented in Dempster and Hutton (1997a, 1999) and Dempster, Hutton, and Richards (1998). The problem is a classical optimal stopping problem that may be formulated as a free-boundary problem by considering the domain partition of optimal stopping. Removing any explicit reference to the free boundary, the option value may be seen to be the unique solution of an order complementarity problem (Borwein and Dempster 1989) by considering its equivalent formulation as a variational inequality and utilizing standard results for coercive operators. Finally, the value is the solution of an abstract linear program (for which a simpler proof than previously is given) which can be solved with standard LP techniques upon suitable domain

Address correspondence to the author at the Centre for Financial Research, Judge Institute of Management Studies, University of Cambridge, Cambridge CB2 1AG, United Kingdom (01223) 339700; e-mail: m.dempster@jims.cam.ac.uk & d.richards@jims.cam.ac.uk; www-cfr.jims.cam.ac.uk.

truncation and discretization. In Section 2 basic numerical methods and our variable coefficient tridiagonal simplex algorithm, together with the degenerate PDE approach (Wilmott, Howison, and Dewynne 1993; Barraquand and Pudet 1996) to valuing market-traded, discretely sampled exotic options, are also reviewed.

Section 3 discusses both theory and basic numerical methods for pricing exotic options with the local volatility surface implied by market values of European options. This is an area pioneered in Rubinstein (1994) and Dupire (1997) and, although some new theoretical proofs are given in Dempster and Richards (1999) and reliable numerical methods are developed in this paper, its treatment here is mainly seen as a vehicle to demonstrate the generality and efficiency of the LP valuation algorithm in Section 2. In Section 4 results for FTSE 100 exotic American index options—fixed-strike Asian puts—are presented to substantiate these claims. First, representative numerical results for vanilla and constant volatility path-dependent American options are presented. Then vanilla European and American options fitting the smile are studied in order to evaluate potential pricing errors in fitting the local volatility surface used to finally price American exotics. Conclusions are drawn in Section 5 and some directions of current and future work are indicated.

2. PRICING AMERICAN OPTIONS USING LINEAR PROGRAMMING

Consider the standard (Black and Scholes 1973) economy, in which there are two financial instruments: a "risky" asset with price $\mathbf{S}$ modeled by a geometric Brownian motion (GBM) and a savings account whose balance is continuously compounded at a constant *risk-free rate* $r \geq 0$. Define an *equivalent martingale* (or *risk neutral*) probability *measure* (EMM) $\mathbb{Q}$ (see Harrison and Kreps (1979) and Harrison and Pliska (1981)) 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 $d\mathbf{S} = r\mathbf{S}\,dt + \sigma\mathbf{S}\,d\tilde{\mathbf{W}}$, where $t \in [0, T]$, $S(0) > 0$, $\sigma > 0$ is the constant volatility of the stock price, and $\tilde{\mathbf{W}}$ is a Wiener process under $\mathbb{Q}$.

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 fixed strike price K exactly 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 these assumptions an American call option will optimally be held to maturity, we shall formulate a version of the American put problem which is suitable for numerical solution. Define the value function $v\colon \mathbb{R}^+ \times [0, T] \to \mathbb{R}$ 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 as $\psi(\mathbf{S}(\tau)) := (K - \mathbf{S}(\tau))^+$ received by the holder upon exercise at a general *stopping time* $\tau \in [0, T]$.

The value function of an American put option can be considered to be the solution to a classical optimal stopping problem—choose the stopping time $\rho(t)$ that maximizes the conditional expectation under $\mathbb{Q}$ of the discounted payoff—and may be shown to be the first time the value falls to the payoff at exercise, that is, $\rho(t) := \inf\{s \in [t, T] : v(\mathbf{S}(s), s) = \psi(\mathbf{S}(s))\}$. 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. Hence $\mathcal{C} := \{(x, t) \in \mathbb{R}^+ \times [0, T) : v(x, t) > \psi(x)\}$ and $\mathcal{S} := \{(x, t) \in \mathbb{R}^+ \times [0, T) : v(x, t) = \psi(x)\}$.

On the continuation region, the value function satisfies the Black–Scholes parabolic partial differential equation (PDE) $\mathcal{L}_{BS}v + (\partial v/\partial t) = 0$ for $(x, t) \in \mathbb{R}^+ \times [0, T]$, where the elliptic operator $\mathcal{L}_{BS} := \frac{1}{2}\sigma^2 x^2(\partial^2/\partial x^2) + rx(\partial/\partial x) - r$, since the discounted stopped price process of the option is a martingale, while as soon as the process crosses into $\mathcal{S}$, $v = \psi$, and to preclude arbitrage $\mathcal{L}_{BS}v + (\partial v/\partial t) \leq 0$. Hence we have

$$\left(-\mathcal{L}_{BS}v - \frac{\partial v}{\partial t}\right) \wedge (v - \psi) = 0 \tag{2.1}$$

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 remove any reference to the optimal stopping boundary by formulating the problem in terms of (2.1) 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 $-\mathcal{L}v - (\partial v/\partial t)$ with constant coefficient elliptic part $\mathcal{L} := \frac{1}{2}\sigma^2(\partial^2/\partial\xi^2) + \left(r - \frac{1}{2}\sigma^2\right)(\partial/\partial\xi) - r$ 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 $\psi(\xi) := \left(K - e^{\xi}\right)$. It will also be convenient to reverse time—to *remaining time to maturity*, for which we will again use the symbol t for simplicity—so that the payoff function ψ becomes an *initial* condition for the Black–Scholes PDE. The American put value function is then the unique solution to

$$\text{(OCP)} \quad \begin{cases} v(\cdot, 0) = \psi \\ v \geq \psi \\ -\mathcal{L}v + \frac{\partial v}{\partial t} \geq 0 \\ \left(-\mathcal{L}v + \frac{\partial v}{\partial t}\right) \wedge (v - \psi) = 0 \quad \text{a.e. in } \mathbb{R} \times [0, T] \end{cases} \tag{2.2}$$

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 (Cryer and Dempster 1980; Borwein and Dempster 1989). Dempster and Hutton (1999) (see also Jaillet, Lamberton, and Lapeyre (1990)) 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 $-\mathcal{L}$ is *coercive*; that is, $\exists \alpha \in \mathbb{R}^+$ 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 program given by

$$\text{(LP)} \quad \inf_u \langle u, c \rangle \text{ s.t. } u \in F \quad \text{for any } c > 0 \text{ a.e. on } \mathbb{R} \times [0, T], \tag{2.3}$$

where

$$F := \left\{ u \in H : u(\cdot, 0) = \psi,\ u \geq \psi,\ -\mathcal{L}u + \frac{\partial u}{\partial t} \geq 0 \right\} \subset P \tag{2.4}$$

since the linear operator $-\mathcal{L}$ on the Hilbert space H is of *type-Z*; that is, $v \wedge y = 0$ ($\Rightarrow \langle v, y \rangle = 0) \Rightarrow \langle v, -\mathcal{L}y \rangle \leq 0 \ \forall v, y \in H$. This basic result giving equivalence between

(OCP) and (LP) is an extension to the parabolic case of a result of Cryer and Dempster (1980) for elliptic partial differential operators. To see this we shall need some notation.

Define the Sobolev space $W^{m,\,p,\,\mu}(\mathbb{R}_2)$ as the space of functions $u \in L^p(\mathbb{R}_2, e^{-\mu|x|}\,dx)$ whose *weak derivatives* of order not exceeding $m \in \mathbb{N}$ exist and are also in $L^p(\mathbb{R}_2, e^{-\mu|x|}\,dx)$, for $p \in [0, \infty]$ and $\mu \in (0, \infty)$. (Here $|\cdot|$ denotes the L^1 norm on $\mathbb{R}^2$ and dx denotes the Lebesgue measure on $\mathbb{R}_2$, and it should be noted that the extension of the results in the sequel to $\mathbb{R}_{n+1}$, for arbitrary $n \in \mathbb{N}$, is completely straightforward.) Consider the Hilbert space $H^1(\mathbb{R}_2) := W^{1,\,2,\,\mu}(\mathbb{R}_2)$, for some fixed $\mu > 0$, of square integrable functions with square integrable derivatives defined on $\mathbb{R}_2$. The Hilbert space $H^1(\mathbb{R}_2)$ has as Banach dual the Sobolev space $H^{-1}(\mathbb{R}_2) := W^{-1,\,2,\,\mu}(\mathbb{R}_2)$, also a Hilbert space of Radon measures, with which it may be identified. Consider the *pairing* $\langle\cdot,\cdot\rangle : H^1 \times H^{-1} \to \mathbb{R}$ between dual spaces given by

$$\langle u, v\rangle := \int_{\mathbb{R}_2} u(\xi, t)v(\xi, t)e^{-\mu(|\xi|+|t|)}\,d\xi\,dt, \tag{2.5}$$

where we may interpret $v \in H^{-1}$ as the density function of the Radon measure element of the dual space H^{-1} of H^1 with respect to $e^{-\mu(|\xi|+|t|)}\,d\xi\,dt$. Alternatively, we may consider $\langle\cdot,\cdot\rangle$ given by (2.5) as an inner product on the Hilbert space $H^0(\mathbb{R}_2) := L^2(\mathbb{R}_2, e^{-\mu|x|}dx)$ by virtue of the canonical injections $H^1 \hookrightarrow H^0 \hookrightarrow H^{-1}$, see Baiocchi and Capelo (1984, p. 79). In this setting the partial differential operator $\mathcal{L}$ may be interpreted either as a map $H^1 \to H^{-1}$ or as an operator on H^1. Consider also the *bilinear form* $a(\cdot,\cdot) : H^1 \times H^1 \to \mathbb{R}$ given by

$$a(u, v) := \int_{\mathbb{R}_2} \left(\frac{\sigma^2}{2}u_\xi v_\xi - \left(\left(r - \frac{\sigma^2}{2}\right) + \mu\frac{\sigma^2}{2}\frac{\xi}{|\xi|}\right)u_\xi v + ruv\right) \times e^{-\mu(|\xi|+|t|)}\,d\xi\,dt. \tag{2.6}$$

which can be chosen to satisfy $a(u, v) = \langle v, -\mathcal{L}u\rangle$, $u, v \in H^1$.

Finally, note that H^1 (and hence H^{-1}) is a vector lattice Hilbert space (but *not* a Hilbert lattice) with positive cone defined in terms of (Lebesgue) almost everywhere nonnegativity. See Cryer and Dempster (1980); Baiocchi and Capelo (1984) and Borwein and Dempster (1989, p. 553–554) for more details on these ideas, which have been adapted here to match the more general setting of Jaillet et al. (1990). In particular, we shall assume all functions in H^1 ($\cong H^{-1}$) considered to be defined as $u(\cdot, |t|)$ on $\mathbb{R} \times (-\infty, 0)$ and as $u(\cdot, T)$ on $\mathbb{R} \times [T, \infty)$ (see Cryer and Dempster (1980, pp. 89 ff)).

The following results (Cryer and Dempster 1980; Dempster and Hutton 1999) relate the bilinear form $a(\cdot,\cdot)$ to the elliptic part of the partial differential operator $\mathcal{L}$ and show variational inequality (VI) and (OCP) equivalence.

THEOREM 2.1. *The variational inequality (VI) given by*

$$\text{(VI)}\qquad \begin{cases} v(\cdot, T) = \psi \\ v \geq \psi \\ u \geq \psi \quad \text{a.e.} \Rightarrow a(v, u - v) + \left\langle u - v, \frac{\partial v}{\partial t}\right\rangle \geq 0 \quad \text{a.e. in } [0, T] \end{cases} \tag{2.7}$$

is equivalent to the order complementarity problem (OCP) and is uniquely solvable.

It can be shown that a given by (2.6), and hence $-\mathcal{L}$, is coercive (see Jaillet et al. (1990, p. 267), whose spaces $L^2([0,T], V_\mu)$ and $L^2([0,T], H_\mu)$ may be considered restrictions respectively of our spaces H^1 and H^0). Then the Lions–Stampacchia theorem (Baiocchi and Capelo 1984) implies that the solution to (VI) is unique. The formulation (VI) is a type of classical physical problem, termed the *(Stefan) obstacle problem*, where the payoff function ψ is the obstacle below which the solution cannot fall.

Next define, for a closed subset $F \subseteq P \subseteq H$ of a vector lattice Hilbert space H with positive cone $P := \{u \in H : u \geq 0\}$, the least element problem

$$\text{(LE)} \qquad v \in F \quad \text{s.t. } v \leq u, \ \ u \in F.$$

The least element v is denoted by $\mathrm{LE}(F)$. Note that if it exists, the least element is always unique since, if v_1 and v_2 are least elements of F, then $v_1 \leq v_2$ and $v_2 \leq v_1$, so from the vector lattice property $v_1 = v_2$.

THEOREM 2.2. *In the setting described above, if $\mathcal{T}$ is a coercive type Z temporally homogeneous elliptic differential operator and F is given by (2.4), then there exists a unique solution v to the following equivalent problems:*

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

$$\text{(LE)} \qquad \textit{find } v = \mathrm{LE}(F),$$

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

Proof. We first prove the equivalence between (OCP) and (LE), after making the trivial domain extensions of the problem functions given above to set them in H^1. Let L denote the Laplace transform operator with respect to the measure $e^{-\mu|t|}$, so that, for $(\xi, \lambda) \in \mathbb{R}_2$, the Laplace transform $\hat{u} \in H^1$ of a function $v \in H^1$ is defined by

$$\hat{u}(\xi,\lambda) := Lu(\xi,\cdot)(\lambda) := \int_0^\infty e^{-|\lambda| t} u(\xi,t) e^{-\mu t} dt. \tag{2.8}$$

As noted above, we have extended the temporal domain of our value functions v to $[0,\infty)$ as constant on (T,∞), so that this generalized Laplace transform is well defined. L is a linear operator and $\mathcal{T}$ is temporally homogeneous (i.e., it has time-independent coefficients), and therefore commutes with the Laplace operator, so that taking the Laplace transform of the operator $\mathcal{T} + \frac{\partial}{\partial t}$ gives $\mathcal{T}L + L\frac{\partial}{\partial t}$. The Laplace transform of the first-order time derivative is given by

$$\begin{aligned} \left(L\frac{\partial v}{\partial t}\right)(\xi,\lambda) &:= \int_0^\infty e^{-|\lambda| t} \frac{\partial v}{\partial t}(\xi,t) e^{-\mu t}\, dt \\ &= -v(\xi,0) + (|\lambda| + \mu)\hat{v}(\xi,\lambda) \end{aligned} \tag{2.9}$$

and $v(\xi,0)$ is given by the initial condition $v(\cdot,0) = \psi$ (in backwards time).

Now, note that the Laplace transform is *positivity-preserving* in the sense that $v \geq 0 \Rightarrow \hat{v} \geq 0$ a.e. on $\mathbb{R}_2$. Then, writing the initial condition, constant in λ, as $\hat{q}(\cdot, \lambda) :\equiv -v(\cdot, 0)$ to agree with the notation of Borwein and Dempster (1989), (OCP) is equivalent to the transformed order complementarity problem $\widehat{\text{(OCP)}}$, also posed in H^1, given by

$$\widehat{\text{(OCP)}} \qquad \begin{cases} \hat{v} \geq \hat{\psi} \\ \mathcal{T}\hat{v} + \hat{q} \geq 0 \\ (\mathcal{T}\hat{v} + \hat{q}) \wedge (\hat{v} - \hat{\psi}) = 0 \quad \text{a.e. on } \mathbb{R}_2, \end{cases}$$

where $\hat{\mathcal{T}} := \mathcal{T} + |\lambda| + \mu$ and $\hat{\psi}$ is the Laplace transform of the log-transformed payoff function ψ, given by $\hat{\psi}(\xi, \lambda) = \psi(\xi)/(|\lambda| + \mu)$. $\hat{\mathcal{T}}$ remains coercive and it is easy to check from the definition that it remains type Z. Consider solutions to the projected (OCP) obtained from $\widehat{\text{(OCP)}}$ by fixing $\lambda \in \mathbb{R}$. We can now apply the order complementarity–least element equivalence result of Borwein and Dempster for coercive type Z *elliptic* operators, so that for each $\lambda \in \mathbb{R}$ one such, $\hat{v}(\cdot, \lambda)$, is the solution (necessarily unique) to the least element problem defined by LE($\hat{F}_\lambda$), where $\hat{F}_\lambda$ is defined by

$$\hat{F}_\lambda := \left\{ \hat{v}(\cdot, \lambda) : \hat{v}(\cdot, \lambda) \geq \psi(\cdot, \lambda), \mathcal{T}\hat{v}(\cdot, \lambda) + \hat{q}(\cdot, \lambda) \geq 0 \right\}. \tag{2.10}$$

It follows that $\hat{v}$ is the unique solution to the least element problem $\widehat{\text{(LE)}}$ defined by $LE(\hat{F})$, where $\hat{F}$ is defined by $\hat{F} := \{\hat{v} : \hat{v} \geq \hat{\psi},\ \mathcal{T}\hat{v} + \hat{q} \geq 0\}$. Applying the inverse Laplace transform L^{-1} to $\hat{v}$ shows that $v = L^{-1}\hat{v}$ solves both (LE), given by LE(F), and (OCP), as required. Indeed suppose the contrary: that there exists $u \in F$ such that $u \leq v$, $u \neq v$. Then it follows since L is positivity preserving that $\hat{u} \in \hat{F}$ and $\hat{u} \leq \hat{v}$, $\hat{u} \neq \hat{v}$, a contradiction to $\hat{v} = \text{LE}(\hat{F})$.

With this least element result, the LP equivalence is immediate: v is the least element of $F \iff v \leq u$ for all $u \in F$, and so $\langle c, v \rangle \leq \langle c, u \rangle$ for all $u \in F$ and any vector $c > 0$. Therefore u minimizes $\langle c, v \rangle$ over all v in F and is thus the solution to the abstract linear program (LP). Restricting to the original problem domain yields the result. □

It should be noted that the above proof depends on time running “backwards” since otherwise we cannot substitute ψ for $v(\cdot, 0)$ in (2.9). The least element result tells us that the linear constraint set lies within the positive cone translated so that its apex lies at v. We can pick out the least element of the constraint set by minimizing $\langle c, u \rangle$ over the set $u \in F$, where $c > 0$; specifically in $\mathbb{R}^2$ by minimizing the intercept of negatively sloped lines defined by $c'u$ with normal $c > 0$ intersecting F.

Theorem 2.2 gives equivalence between (VI), (OCP), (LE), and (LP) for the American put, since $-\mathcal{L}$ is coercive type Z (see Jaillet et al. 1990). It should be stressed that the result is very general and applies to virtually any parabolic partial differential operator with a *temporally homogeneous* coercive type Z elliptic part, and virtually any payoff function. For example, it may be applied to the Black–Scholes operator $-\mathcal{L}_{BS}$ directly without prior logarithmic transformation of the space variable. A more delicate argument involving step function coefficient approximation and a suitable passage to the limit can be used to establish the results of Theorem 2.2 for *time-dependent* coefficient operators (the details will appear elsewhere) which are required in this paper to handle the time-varying nature of local volatility coefficients. Theorem 2.2 also suggests a simple way to solve the equivalent problems numerically: by a suitable discretization the infinite-dimensional abstract linear program (LP) reduces to an ordinary linear program with solutions in $\mathbb{R}^n$. Thus we next discretize the problem and consider efficient LP numerical solution methods.

2.1. Computation

We shall approximate the value function by a function that is piecewise constant on rectangular intervals between points in a regular lattice. Approximating the partial derivatives by standard Crank–Nicolson finite differences (see, e.g., Wilmott et al. 1993) we obtain a discrete form of (OCP) that can be rewritten in matrix form upon collapsing the space index. A matrix is type-Z if it has nonnegative off-diagonal elements (see, e.g., Borwein and Dempster 1989), which in the case of the matrix derived from the discretized negative Black–Scholes operator 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 (Dempster and Hutton 1999). From this condition it can also be shown that this matrix is coercive (Jaillet et al. 1990; Hutton 1995).

The LP formulation can be solved in backwards time either directly or iteratively and the interested reader can find comparisons of solution methods in Hutton (1995) and Dempster and Hutton (1999). Here we solve the discretized (LP) using time-stepping and a simplified revised simplex algorithm that takes advantage of the tridiagonal structure of the constraint matrix formed from standard Crank–Nicolson finite difference approximations to produce a fast accurate direct solution method detailed in Dempster et al. (1998) and Richards (1999). This procedure is suitable for any standard constant parameter Black–Scholes type formulation, but also yields significant computational savings for valuation problems with volatility and drift parameters which are functions of time. It incorporates a technique for the solution of problems with nonconstant 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 Dempster et al. (1998) results are presented for this updating procedure which show that even for a general constraint matrix the procedure outperforms standard commercial LP solvers by orders of magnitude.

To understand the exact computational savings of these simplex methods, first consider the complexity of the vanilla 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 simplex basis is $\mathcal{O}(I)$ and hence we have $\mathcal{O}(I)$ 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 4.3 later)—at most one new basic variable enters at each time step. Far from maturity, calculations for several time steps may even 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 $\mathcal{O}(MI)$, where M is the number of time steps.

For the basis factorization updating technique required by each simplex iteration (space step) the calculations result in a similar complexity, but can be performed in three floating point operations, although extra computation time is needed for the dynamic allocation of the upper-lower (UL) factorization. Results for the constant coefficient method and for the nonconstant coefficient updating technique are reported in Dempster et al. (1998), along with results for a complete calculation of the full LU factorization at each iteration to highlight the overheads of using general commercial solvers.

2.2. Extension to Exotic Options

An *exotic option* is any derivative security that has a path-dependent component in its payoff at exercise. We may formulate discretely sampled exotic option valuation problems as linear programs through state augmentation, particularly for American discretely sampled lookback and Asian options. Exotic option values are dependent on the underlying stock price, (forward) time, and an additional "independent" variable that encapsulates the required path information.

We now outline the formulation of a generic American exotic option in a discretely sampled setting using the unifying framework of Wilmott et al. (1993). 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 value of a path-dependent variable, such as the *average* in the case of Asian options, or the *maximum/minimum* for lookback options. Assuming that the asset price is sampled on N occasions during the life of the option with maturity T, denote by M_n the observed value of the augmented variable at the sampling date t_n, $n = 0, \ldots, N-1$, so that sampling begins at time 0 and $t_0 = 0$. For completeness, define $t_N := T$. The variable M_n is a constant value throughout the period $[t_n, t_{n+1})$, since no sampling takes place until time t_{n+1}. Effectively M_n is 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 Wilmott et al. 1993 for more details). Across a general sampling date t_n the augmented variable is updated from a value M_{n-1} just prior to the date to a value M_n at the sample date. No-arbitrage arguments lead to the jump condition $V\left(S, M_{n-1}, t_n^-\right) = V\left(S, M_n, t_n\right)$ $n = 1, \ldots, N-1$, where t_n^- is a time immediately before the sampling date t_n.

In the time interval $[t_n, t_{n+1})$ the European value V satisfies the augmented Black–Scholes PDE defined by $\mathcal{L}_{BS} + f(S, t)\partial V/\partial M + 2V/2t$, where $f(S, t)$ is a function specified for the option of interest. We consider the final period $[t_{N-1}, T]$ and use a dynamic programming algorithm to determine values for earlier periods. As in the vanilla American put case treated above, the American exotic valuation domain in $\mathbb{R}^3$ 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 (Dempster and Richards 1999). To complete the formulation of the discretely sampled exotic value in the final period we require a terminal condition $V(S, M_{N-1}, T) = \psi(S, M_{N-1})$ for all $S, M_{N-1} \in \mathbb{R}^+$ and boundary conditions in S at $S = 0$ and as $S \to \infty$ which are option dependent and are discussed in more detail in the cited paper. If 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 $\mathcal{L}$, we may define a new partition with regions $\mathcal{C}_N$ and $\mathcal{S}_N$. 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 of the form (2.2) over $\mathbb{R}^2 \times (t_{N-1}, T]$ involving payoff $\psi(\xi, \zeta_{N-1}) := \max\left(e^{\zeta_{N-1}} - e^{\xi}, 0\right)$ with V also denoting the option value as a function of ξ and ζ_{N-1}. This puts us in a framework equivalent to that of the vanilla American put of Section 2 but with the additional parameter ζ_{N-1}, and hence we have equivalence to an abstract LP for each value of the augmented variable $\zeta_{N-1} \in (-\infty, \infty)$. This problem must be solved for all possible values of the parameter ζ_{N-1}. Applying the jump conditions 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. THE INVERSE PROBLEM OF OPTION PRICING

The Black–Scholes economy has one unobservable quantity—the volatility parameter σ—which must be inferred. Particularly since the stock market crash of 1987 the volatility of equity options has exhibited variation both with the strike price—the *volatility smile* (or *skew*)—and the option's maturity—the volatility *term structure*. These effects highlight the market's deviation from the Black–Scholes assumption that the future asset price has a constant variance lognormal conditional probability density.

Several approaches have been suggested in the literature to model this behavior. One approach is to treat the volatility as a second stochastic factor with the aim of specifying the time-varying volatility from the model (Hull and White 1987; Clarke and Parrott 1998). However this approach is difficult to fit to market data, is not arbitrage free, and introduces an additional dimension to the pricing problem. An alternative one-factor approach allows the volatility $\sigma := \sigma(S, t)$ to be a variable that is both state and time dependent. By starting from the market data and finding the local volatilities that are consistent with the market—commonly termed an *inverse problem* (Andersen and Brotherton-Ratcliffe 1997; Lagnado and Osher 1997; Bouchouev and Isakov 1999)—this model can be made to price the market nearly exactly. 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 (Rubinstein 1994; Derman and Kani 1998). All methods used to fit market data are prone, however, 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 and allow arbitrage possibilities in the 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. Therefore, it is often assumed that a *continuum* of European call option prices $C(K, T)$ are available for all strikes and maturities. Alternatively, recent papers have implemented *regularization* methods (Jackwerth and Rubinstein 1996; Bodurtha 1997; Lagnado and Osher 1997; Coleman, Li, and Verma 1999) to make the inverse problem stable.

All approaches to fitting the smile suffer from the consequences of inconsistent data and will not price correctly *all* options—particularly those far out of the money—in the face of these data problems. The modeling approach used in this paper is not claimed to be the most accurate or efficient, but it is quick and highlights the versatility of the LP pricing method in the face of a degenerate ill-posed problem with nonconstant coefficients.

3.1. 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 Dupire (1997) and has since been given a more formal treatment in Derman and Kani (1998). The main idea is that there exists an *adjoint* or *dual* PDE to the Black–Scholes PDE with strike price K and maturity T as the independent variables, which can be derived through consideration of the conditional probability distribution of the underlying stochastic process and the forward PDE satisfied by this probability density.

We assume that a continuum of European call option prices $C(K, T)$ for all strikes and maturities $K, T \in \mathbb{R}^+$ are available from the market and any gaps in the data can be filled by interpolation or extrapolation techniques. We will deal with any arbitrage violations in these approximated values later. The underlying asset price is assumed to follow a diffusion process under the risk-neutral measure $\mathbb{Q}$, but now with *nonconstant* volatility; that is, $d\mathbf{S} = \mathbf{S}r(t)\,dt + \mathbf{S}\sigma(\mathbf{S}, t)\,d\mathbf{W}$. The price of a European call option can be written in terms of an expectation under $\mathbb{Q}$ with respect to the conditional probability density function (pdf) $p(s, T|S, t)$ of the underlying asset $\mathbf{S}$ having value s at time T given that the asset price at time t is S as

$$C(S, t; K, T) = P(t, T) \int_0^\infty p(s, T|S, t)(s - K)^+ \, ds, \tag{3.11}$$

where $t \leq T$ and the discount factor $P(t, T) := \exp\left(-\int_t^T r(u)\, du\right)$.

Breeden and Litzenberger (1978) showed that the European call option price and this conditional probability density were related as $p(K, T|S, t) = P(t, T)^{-1} \partial^2 C(S, t; K, T)/\partial K^2$ by differentiation of (3.11). The function $p(K, T|S, t)$ is also the Green's function (or *fundamental solution*) of the Black–Scholes PDE. Thus it satisfies this 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 idealized market data. The function $p(K, T|S, t)$ can also be shown to satisfy the Fokker–Planck (or *forward Kolmogorov*) PDE utilizing the following theorem (see, e.g., Jackwerth and Rubinstein 1996).

THEOREM 3.1. *The conditional pdf* $p(y, \tau|x, t)$ *of a general stochastic process* $\mathbf{X}(t)$ *where* $t \geq 0$ *given by* $d\mathbf{X}(t) = \mu(\mathbf{X}, t)\, dt + \sigma(\mathbf{X}, t)\, d\mathbf{W}(t)$ *satisfies the Fokker–Planck or forward Kolmogorov equation*

$$\frac{\partial p(y, \tau|x, t)}{\partial \tau} + \frac{\partial\, (\mu(y, \tau) p(y, \tau|x, t))}{\partial y} - \frac{1}{2} \frac{\partial^2 \left(\sigma^2(y, \tau) p(y, \tau|x, t)\right)}{\partial y^2} = 0, \tag{3.12}$$

for fixed $(x, t) \in \mathbb{R} \times \mathbb{R}^+$ *with initial condition* $p(y, t|x, t) = \delta(x - y)$.

COROLLARY 3.2. *The transitional probability density function* $p(S', T|S, t)$ *of the stock price process* $d\mathbf{S}/\mathbf{S} = r(t)\, dt + \sigma(\mathbf{S}, t)\, d\mathbf{W}$ *satisfies the PDE*

$$\frac{\partial p(S', T|S, t)}{\partial T} = \frac{1}{2} \frac{\partial^2 (\sigma^2(S', T) S'^2 p(S', T|S, t)}{\partial S'^2} - \frac{\partial}{\partial S'} \left(r S' p(S', T|S, t)\right) \tag{3.13}$$

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

The following corollary is due to Dupire (1997).

COROLLARY 3.3. *Given that the underlying asset price process in Corollary* 3.2 *the price of a European call option* $C(S, t; K, T)$ *solves the partial differential equation*

$$\frac{\partial C}{\partial T} = \frac{\sigma(K, T)^2 K^2}{2} \frac{\partial^2 C}{\partial K^2} - r(T) K \frac{\partial C}{\partial K} \tag{3.14}$$

with boundary condition $C(S, t; S, t) = 0$ (*see also Derman and Kani 1998*).

3.2. Computation

The local volatility function $\sigma(S, t)$ can be fully determined from (3.14) since all other terms in the equation can be found from market data. However, since the inverse problem is ill-posed there may not be a unique local volatility functional that fits the market data. To obtain an approximately arbitrage-free pricing algorithm the implied volatility (or, alternatively, call price) data available from the market must be fitted exactly. To this end we apply cubic splines to approximate the implied volatility surface. The ease with which approximations to first and second derivatives can be obtained from a spline fit to market data is a major advantage in reducing the computational complexity of our method since the construction of a length N cubic spline is the solution of a tridiagonal system and hence an $\mathcal{O}(N)$ calculation. The complexity of the problem can be further reduced by *precomputing* derivative information, but at the expense of additional memory requirements.

We outline now the approach used to obtain a consistent *local* volatility surface that fits the market-implied volatility data for European call options. This will be used in the linear programming approach of Section 2 to price American exotic options consistent with the volatility smile. Since we consider the implied volatility nonconstant in the underlying diffusion, assume that the quoted European call prices are Black–Scholes prices with implied volatilities as an additional parameter and define the call prices in terms of the implied volatility $\nu(K,T)$ for the call option of strike K and maturity T as $C(K,T) := C_{BS}(S,t;\, K,T,\nu(K,T))$, 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 as follows.[1]

THEOREM 3.4. *Let the asset price* **S** *follow the diffusion process in Corollary* 3.2. *Then the local volatility function $\sigma(S,t)$ consistent with the arbitrage-free European call prices is given uniquely, in the absence of dividends, by*

$$\sigma^2(K,T) = 2\left(\frac{\partial C}{\partial T} + r(T)K\frac{\partial C}{\partial K}\right)\left(K^2\frac{\partial^2 C}{\partial K^2}\right)^{-1} \tag{3.15}$$

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

$$\sigma^2(K,T) = \frac{2\left(\gamma\frac{\partial\nu}{\partial T} + \frac{1}{2}\frac{\nu}{\gamma} + r(T)K\frac{\partial\nu}{\partial K}\right)}{K^2\left(\frac{1}{\nu}\left(\frac{\partial\nu}{\partial K}\left((T-t)-d_1\right) - \frac{1}{K\gamma}\right)\left(-d_1\frac{\partial\nu}{\partial K}\gamma - \frac{1}{K}\right) + \frac{\partial^2\nu}{\partial K^2}\gamma + \frac{\gamma}{K}\frac{\partial\nu}{\partial K}\right)}, \tag{3.16}$$

where

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

Proof. Equation (3.15) follows immediately from (3.14). By portfolio dominance arguments (see Andersen and Brotherton-Ratcliffe 1997) it can be shown that $\sigma^2(K,T)$ is nonnegative in the absence of arbitrage if the numerator of (3.15) is nonnegative. The denominator is the Breeden–Litzenberger conditional (transitional) probability density which must be nonnegative. Equation (3.16) follows from (3.15) after much calculation (see Richards 1999). □

We will use (3.16) to obtain the local volatilities. For our method the implied volatilities are relatively stable and the expression for the local volatility involves relatively few

[1] This result was derived following discussions with S. H. Babbs (of Bank One) and has since been independently presented in Andersen and Brotherton-Ratcliffe (1997).

numerical calculations and only one computationally expensive logarithmic calculation for each evaluation point. Splines are fitted to the market implied volatility data, with the calculated second derivatives with respect to the maturity stored in an array. By fitting the splines across maturities for each strike first we obtain approximations for the first and second derivatives with respect to the strike. However the first-order derivative of the volatility with respect to maturity is not a natural by-product of the interpolation and is specified here 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 (3.16).

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. At each node (i, m) of the transformed 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 in Section 2 except that the volatility in the formulation is no longer assumed constant. This changes the constraint matrix in the order complementarity problem, and thus in the linear programming formulation; the matrix now has nonconstant diagonals but is still tridiagonal in nature. However, for American and exotic options this is the ideal problem to be solved by the nonconstant tridiagonal simplex method described in Section 2.

Several studies have aimed at detailing the instabilities in market 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. We filter the data using methods derived from the underlying theory. When upon discretization a negative denominator occurs in (3.16) we necessarily have an arbitrage opportunity appearing since this implies a negative value of the transitional probability density. We correct this value using put-call parity to consider the prices of European put options implied by the call option market data. Indeed, if the market data imply a negative value of $p(K_j)$ for some $j = 0, \ldots, N$, where $p(\cdot) := p(\cdot, T|S, t)$ now denotes the discrete form of the conditional density, we consider a European put option of strike K_{j+1} with value $P(S, t; K_{j+1}, T)$. Then

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

which given that the $P(S, t; K_i, T)$ are known for all $i \leq j$ (from put-call parity) allows us to find a value of the discrete probability $p(K_j)$ consistent with the market data. If this probability is also negative, its value is set to zero. The other possible inconsistency in the data occurs when the numerator of (3.16) is negative, when we set $\sigma^2(K, T)$ to zero.

4. NUMERICAL RESULTS

In this section we present empirical results for the procedure outlined in Section 3 for pricing American options consistent with the observed market volatility smile, together with benchmark results for their constant volatility equivalents. To facilitate comparisons we use as our underlying implied volatility surfaces market data appearing previously in the literature and we price European, American, and exotic options with respect to it. All solution times quoted are for calculations 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 Dempster et al. (1998) for solution on a Pentium II 400 Mhz PC which gives significant speed-ups over those presented here for most levels of domain discretization.

4.1. Vanilla American Put

Table 4.1 illustrates the savings that the new tridiagonal simplex solver makes over the PSOR algorithm[2] (Cryer 1971). The timings in Table 4.1 are CPU times, including all data initialization for the value at time 0 of an at-the-money American put with parameters $K := 1.0$, $T := 1.0$, $\sigma := 20\%$, and $r := 10\%$. The log stock price was bounded above by $U := 2$ and below by $L := -1$, giving the range in untransformed variables as $[0.37, 7.39]$. The number of time steps M was set at 1000 and the number of space steps I varied.

We see from the table that all of our tridiagonal revised simplex methods are linear in space and give impressive speed-ups over PSOR, with the constant-coefficient method (column 2) ranging from 4–500 times faster. For comparison, results are included for the UL update technique (column 3) and solution times for the tridiagonal solver with recalculation of the whole decomposition at each iteration (column 4). The slowest of our solution methods is faster than PSOR at all but the lower levels of discretization and all our methods are able to price 4 or 5 options per second accurately.

TABLE 4.1
Comparison of Tridiagonal Simplex Solvers with PSOR

	Tridiagonal simplex			
Space steps	Constant coefficients	UL update	Recalculation simplex	PSOR
75	0.02	0.05	0.10	0.07
150	0.05	0.08	0.17	0.13
300	0.10	0.14	0.33	0.27
600	0.19	0.24	0.65	1.25
1200	0.38	0.47	1.26	6.37
2400	0.77	1.00	2.47	37.55
4800	1.61	2.24	5.12	255.06
9600	3.71	5.09	10.77	1856.91

Solution CPU times in seconds for $M := 1000$.

[2] See Dempster and Hutton (1999) for comparison of the commercial IBM Optimization Subroutine Library (OSL) with PSOR, interior-point, and simplex methods.

4.2. Discretely Sampled Lookback Strike Options

Next we compare our results for discretely sampled American lookback strike options against some of those in the literature. A lookback strike option has a payoff similar to the corresponding (put or call) vanilla option, but with the strike price replaced by the maximum or minimum realized asset price. For example, an American lookback strike put might have payoff $\psi(S, M) := \max(M - S, 0)$ at exercise, where M is the maximum asset price over the life of the option until exercise, and S is the asset price at exercise as usual. We concentrate on strike options here rather than on rate options (Wilmott et al. 1993) which are treated below. The valuation of American discretely sampled lookbacks is usually achieved using tree methods, though some closed-form solutions are available for the European continuously sampled case.

Our main source of comparison for our numerical results is the PSOR method utilized by Wilmott et al. (1993) after a similarity transform of the augmented Black–Scholes PDE. The sampling schemes investigated, A, B, and C, correspond approximately to samples every 1, 2, and 3 months respectively. The lookback option is of one-year maturity with $\sigma := 20\%$ and $r := 10\%$ and the option is valued at-the-money with initial stock price $S := 100$. Further details and additional results can be found in Dempster et al. (1998) and Richards (1999). Table 4.2 shows the CPU timings and solutions for the sampling schemes. These results are a good comparison because of the explicit description of the sampling schemes employed, but are quoted in terms of the similarity reduced variable to only one decimal place. We have agreement in the results to the accuracy quoted, with solution times of much less than 6 seconds using the constant-coefficients simplex solver, which far outperforms the PSOR method. The higher-order discretizations for the former are given to show convergence. Further comparison results for discretely sampled Asian strike options may be found in Dempster et al. (1998).

4.3. Fitting the Volatility Smile

To test the numerical procedure described in Section 3 we use real FTSE 100 index volatility values implied from the European call option data described in Duan (1995). The implied volatility data are shown graphically in Figure 4.1 and correspond to FTSE 100 volatility values for 31 March 1995. The initial index level was 3129.5 and data were quoted for eight strike prices and five maturities, although for the last two maturities prices were quoted for different strikes. The data were interpolated to fill the gaps as

TABLE 4.2
Comparison of Lookback Valuation Results for the Tridiagonal Method against the Similarity Transformed Method

	Tridiagonal						
	200 × 300 × 200		400 × 600 × 400		800 × 1200 × 800		Similarity Transform
Scheme	Solution	Time	Solution	Time	Solution	Time	Solution
A	10.532	6.46	10.550	53.88	10.555	477.32	10.5
B	9.445	6.69	9.454	59.05	9.457	508.95	9.5
C	8.114	6.59	8.116	59.32	8.116	489.60	8.1

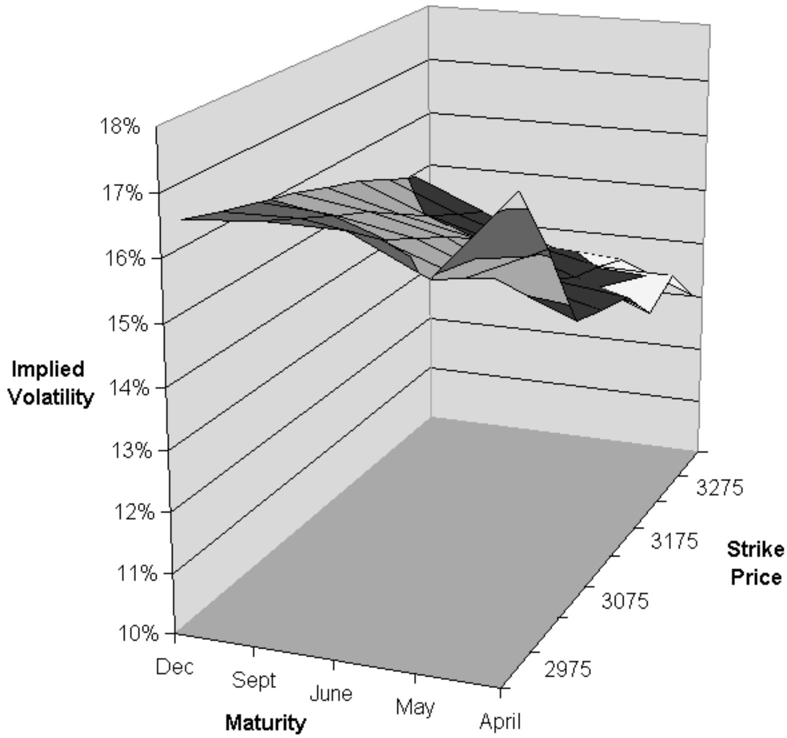

FIGURE 4.1. Implied volatility surface for European call options on the FTSE 100 index.

described in Section 3. We also assume that a constant rate of interest, $r = 10\%$, applies throughout for all maturities.

As described in Section 3 the pricing procedure occurs in two stages. The first—calibrating the local volatility—uses the estimates of the implied volatility derivatives to calculate the surface $\sigma(S, t)$ for use in pricing. In the sequel this calibration takes place on the same mesh that is used in the pricing procedure. There are significant computational savings to be made by calculating the local volatilities on a coarser mesh and performing some type of interpolation between the calibrated values. 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 quote solution times for the option valuation only, assuming that the precomputed local volatility surface is in memory. After calibrating the volatility surface, the options are valued using the UL update algorithm of Section 2 for options with American exercise and a simple linear equation solver for European options. At each time-step for the former, the basis decomposition was calculated and the UL update applied only when new variables entered the basis.

Empirical tests showed that the underlying discretization error due to the Crank–Nicolson numerical approximation of the derivatives was less than 8 basis points on a grid of 3200 time steps and 200 space steps; this must be accounted for in any accuracy comparisons.

4.4. European Option Results

We now use our pricing algorithm to *recover* market option values in order to assess the fit of the calibrated surface. Table 4.3 contains European call option prices for the strikes and maturities given by the market. The reference (actual) values are calculated through the use of the Black–Scholes pricing equation using the FTSE 100 implied volatility values. This requires an approximation for the cumulative normal distribution which is accurate to 10 decimal places. The errors can be seen to be comparable to—and in some cases more accurate than—the baseline numerical accuracy described earlier; in fact all pricing errors are less than 5 basis points at the discretization level used. We conclude that fitting the volatility smile does not induce 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 volatility surface fit to the data is consistent with the local volatilities implied by the market through quoted prices. Solution time for a European option at the discretization employed is approximately 0.6 seconds.

The calculated local volatility surface is displayed in Figure 4.2 and a comparison with the implied volatility surface of Figure 4.1 shows significant differences. The local volatility surface is represented in Figure 4.2 on a truncated strike domain. At short maturities a spike of local volatility occurs at a strike value of approximately 3250, which distorts any graphical representation of the local volatility surface but does not cause instabilities 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 and illustrates that for short maturities the local volatility behaves like the reciprocal of the conditional transitional probability density, smoothing out somewhat for higher maturities.

4.5. Pricing American Options

Next we introduce an added dimension to the problem by pricing American options, since there is a very real possibility that the vanilla optimal exercise boundary (see Figure 4.3) will be moved by fluctuations in the local volatility surface. The 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. The boundary conditions used in this section are as described in Dempster et al. (1998).

TABLE 4.3
Pricing errors $\times 100$ of FTSE 100 European Call Options Fitting the Smile

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

$M = 3200$, $I = 200$. Actual value given by the Black–Scholes call pricing equation.

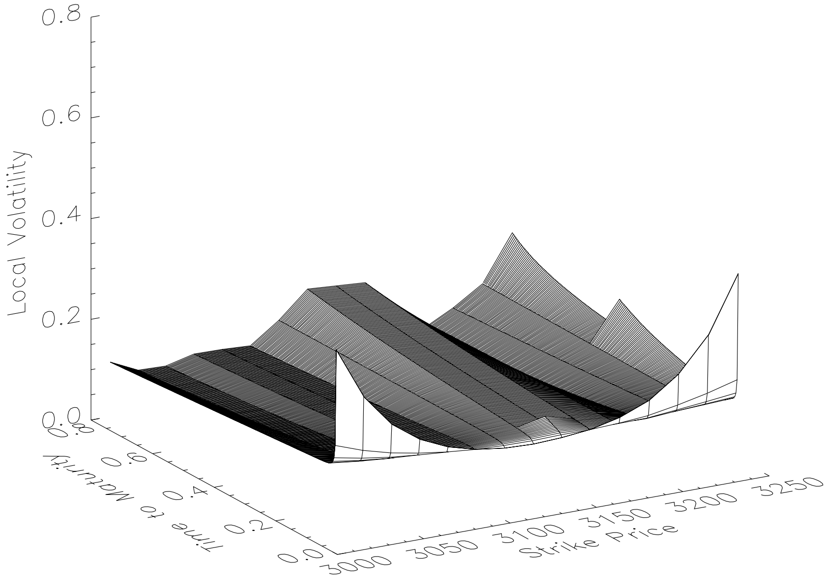

FIGURE 4.2. Local volatilities $\sigma(S, t)$ for FTSE 100 index options on a truncated strike domain. Parameters $M = 200$, $I = 200$ for the full strike range.

The results for the FTSE 100 index in Table 4.4 show American put valuations for the LP tridiagonal solver. As a benchmark we use the original LP valuation on a very fine solution mesh using the at-the-money BS implied volatility. American option values are *higher* than the corresponding constant volatility values to within the numerical tolerance previously evaluated, with the LP solution time being 0.6 seconds. Since both solution algorithms converge to within the same numerical accuracy, the discrepancies between the smile-fitting and at-the-money implied values are due to the volatility surface. It was noted in Section 2 that the convex shape of the optimal exercise boundary for the vanilla American put problem was useful in increasing the efficiency of the tridiagonal solver. Figure 4.3 highlights the reason why the volatility smile fitting option value is different from its vanilla counterpart by illustrating the shifted nonconvex shape of its optimal exercise boundary. When we take account of local volatility, the exercise boundary is no longer a convex function of the asset price, but is shifted horizontally by changes in local volatility. Although this poses no problem for the accuracy of the pricing algorithm, it does radically affect the realized option price. Similar results for the S&P 500 index options are given in Dempster and Richards (1999).

4.6. Pricing American Discretely Sampled Asian Fixed-Strike Options

Finally we price Asian fixed-strike options fitting the smile. An Asian rate (or fixed-strike) option has the vanilla payoff with the asset price at exercise replaced by the average. For example, the Asian rate put has payoff $\psi(S, A) := \max(K - A, 0)$ for the fixed strike K.

Table 4.5 contains option values for the American Asian put option with fixed-strike equal to the initial asset price. All results correspond to options of 0.211 year maturity with the risk-free rate assumed constant at 10 percent. In the table the "Implied value" columns refer to the value found using the LP approach with constant volatility set to the

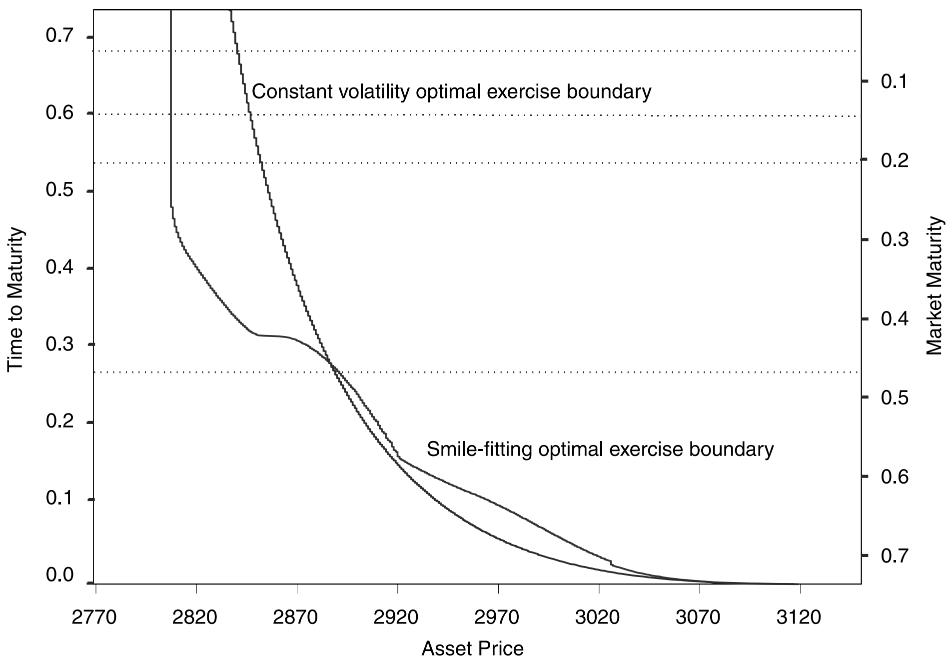

FIGURE 4.3. American put optimal exercise boundary for the smile fit compared with the 1-factor LP method using ATM BS implied volatility. Parameters: $M = 1000$, $I = 3000$, $S_0 = 3129.5$, $K = 3125$, $\sigma_{BS} = 15.89\%$. Dotted lines show market European call maturities.

Black–Scholes at-the-money implied volatility for the options in question. The columns labeled "Smile value" illustrate the results obtained by fitting the volatility smile and term structure.

As can be seen, the option price fitting the volatility smile is significantly different from the constant volatility price. Unlike the vanilla American put the American fixed-strike Asian option results on the FTSE 100 index (Table 4.5) show a mixed effect. For a low arithmetic average sampling rate 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 4.1.

5. CONCLUSION

After surveying earlier results, we applied a fast accurate linear programming valuation algorithm to pricing American exotic 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

TABLE 4.4
American Put Option Valuation Results Fitting the FTSE 100 Volatility Smile

Strike price	April		May		June	
	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

$M = 3200$, $I = 200$. Solution times approximately 0.6 seconds. LP(σ_{atm}) is the value calculated from the tridiagonal solver with discretization $M = 10000$, $I = 10000$.

those with constant volatilities. Finally, we have seen similar effects for American 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 discontinuous option payoffs, including barriers (Dempster et al. 1999), and high-dimensional Bermudan and American fixed income derivatives (see Dempster and Hutton (1997b)).

TABLE 4.5
Discretely Sampled American Fixed-Strike Asian Put Option Results Fitting the FTSE 100 Smile

2 Samples					12 Samples				
M	I	J	Implied value	Smile value	M	I	J	Implied value	Smile value
400	200	200	0.360	0.295	270	200	200	3.134	3.710
800	200	200	0.360	0.296	540	200	200	3.135	3.759
800	400	400	0.353	0.278	540	400	400	3.127	3.737

Parameters: $K = 2975$, $S_0 = 3129.5$ and $T = 0.211$ years. Solution time for $M = 200$, $I = J = 200$ is approximately 18 seconds.

REFERENCES

ANDERSEN, L. B. G., and R. BROTHERTON-RATCLIFFE (1997): The Equity Option Volatility Smile: An Implicit Finite-Difference Approach, *J. Computat. Finance* 1, 5–37.

BAIOCCHI, C., and A. CAPELO (1984): *Variational and Quasivariational Inequalities*. New York: Wiley.

BARRAQUAND, J., and T. PUDET (1996): Pricing of American Path-Dependent Contingent Claims, *Math. Finance* 6, 17–51.

BLACK, F., and M. SCHOLES (1973): The Pricing of Options and Corporate Liabilities, *J. Political Econ.* 81, 637–659.

BODURTHA, J., Jr. (1997): A Linearization-Based Solution to the Ill-Posed Local Volatility Estimation Problem, Working paper, Georgetown University.

BORWEIN, J. M., and M. A. H. DEMPSTER (1989): The Linear Order Complementarity Problem, *Math. Operations Res.* 14, 534–558.

BOUCHOUEV, I., and V. ISAKOV (1999): Uniqueness, Stability and Numerical Methods for the Inverse Problem that Arises in Financial Markets, *Inverse Prob.* 15, R95–R116.

BREEDEN, D., and R. LITZENBERGER (1978): Prices of State-Contingent Claims Implicit in Option Prices, *J. Business* 51, 621–651.

CLARKE, N., and K. PARROTT (1998): Multigrid for American Option Pricing with Stochastic Volatility, Working paper, Mathematical Sciences, University of Greenwich.

COLEMAN, T. F., Y. LI, and A. VERMA (1999): Reconstructing the Unknown Local Volatility Function, *J. Computat. Finance* 2, 77–102.

CRYER, C. W. (1971): The Solution of a Quadratic Programme Using Systematic Overrelaxation, *SIAM J. Control Optim.* 9, 385–392.

CRYER, C. W., and M. A. H. DEMPSTER (1980): Equivalence of Linear Complementarity Problems and Linear Programs in Vector Lattice Hilbert Spaces, *SIAM J. Control Optim.* 18, 76–90.

DEMPSTER, M. A. H., A. ESWARAN, and D. G. RICHARDS (1999): Wavelets and Their Application in Valuing Derivatives by PDE Methods, Working paper, Centre for Financial Research, University of Cambridge.

DEMPSTER, M. A. H., and J. P. HUTTON (1997a): Fast Numerical Valuation of American, Exotic and Complex Options, *Appl. Math. Finance* 4, 1–20.

DEMPSTER, M. A. H., and J. P. HUTTON (1997b): Numerical Valuation of Cross-Currency Swaps and Swaptions; in *Mathematics of Derivative Securities*, M. A. H. Dempster and S. R. Pliska, eds. Cambridge: Cambridge University Press, 473–503.

DEMPSTER, M. A. H., and J. P. HUTTON (1999): Pricing American Stock Options by Linear Programming, *Math. Finance* 9, 229–254.

DEMPSTER, M. A. H., J. P. HUTTON, and D. G. RICHARDS (1998): LP Valuation of Exotic American Options Exploiting Structure, *J. Computat. Finance* 2, 61–84.

DEMPSTER, M. A. H., and D. G. RICHARDS (1999): Pricing Exotic American Options Fitting the Volatility Smile, Working paper 17/99, Judge Institute of Management Studies, University of Cambridge, submitted.

DERMAN, E., and I. KANI (1998): Stochastic Implied Trees: Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility, *Inter. J. Theoret. Appl. Finance* 1, 61–110.

DUAN, J.-C. (1995): The Garch Option Pricing Model, *J. Finance* 5, 13–32.

Dupire, B. (1997): Pricing and Hedging with Smiles; in *Mathematics of Derivative Securities*, M. A. H. Dempster and S. R. Pliska, eds. Cambridge: Cambridge University Press, 103–111.

Harrison, J. M., and D. Kreps (1979): Martingales and Arbitrage in Multi-Period Security Markets, *J. Econ. Theory* 20, 381–408.

Harrison, J. M., and S. R. Pliska (1981): Martingales and Stochastic Integrals in the Theory of Continuous Trading. *Stoch. Process. Appl.* 11, 215–260.

Hull, J., and A. White (1987): The Pricing of Options on Assets with Stochastic Volatilities, *J. Financial Quant. Anal.* 3, 281–300.

Hutton, J. P. (1995): Fast Pricing of Derivative Securities, Ph.D. thesis, Department of Mathematics, University of Essex.

Jackwerth, J. C., and M. Rubinstein (1996): Recovering Probability Distributions from Option Prices, *J. Finance* 51, 62–74.

Jaillet, P., D. Lamberton, and B. Lapeyre (1990): Variational Inequalities and the Pricing of American Options, *Acta Appl. Math.* 21, 263–289.

Lagnado, R., and S. Osher (1997): A Technique for Calibrating Derivative Security Pricing Models: Numerical Solution of an Inverse Problem, *J. Computat. Finance* 1, 13–25.

Richards, D. G. (1999): Pricing American Exotic Options, Ph.D. thesis, Judge Institute of Management Studies, University of Cambridge.

Rubinstein, M. (1994): Implied Binomial Trees, *J. Finance* 49, 771–818.

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