

An important task in option pricing is to infer the local volatility function of the underlying price process from quoted option prices. **Domingo Tavella** and **Wolfgang Klopfer** achieve this by numerical solution of the Fokker Plank equation



Assume a price process $S(t)$ given by:

$$\frac{dS}{S} = \mu dt + \sigma(S, t) dW.$$

The problem here is to compute an approximation to the local volatility function, $\sigma(S, t)$, from quoted option prices. The implied local volatility task consists of mapping the implied Black and Scholes volatility surface, $\sigma_{BS}(K, T)$, where K is the strike and $0 \leq T \leq T_{max}$ is the maturity, onto the local volatility surface $\sigma(S, t)$, where S is the underlying price and $0 \leq t \leq T_{max}$ is local time.

A number of approaches have been proposed to accomplish this objective, varying in numerical complexity and analytical sophistication. Among such methodologies are the use of implied trees, the numerical solution of the pricing equation with suitable approximations for the local volatility function, and the use of relative entropy arguments to get approximations to the local volatility function. The varying degrees of success of these methods has been broadly discussed in the literature. It is well-known, for example, that implied trees may face significant difficulties in successfully implying the local volatility function. This difficulty can be traced to the fact that trees are essentially an explicit finite difference formulation of the pricing problem, with the rigidity characteristic of explicit schemes. The solution of the pricing equation is a more robust approach, but it is also numerically time-consuming. The relative entropy approach is based on minimizing the relative entropy with respect to an assumed prior distribution.

In this article we demonstrate the use of numerical solutions of the Fokker-Plank equation to economically extract the implied local volatility function from quoted prices of options. We illustrate the methodology using European call options for the sake of simplicity. The methodology is general and can be used with any type of options.

It is also well-known that, in principle, the local volatility function can be obtained directly from the pricing equation for the forward price of a European call option. This well-known result, due to Dupire, is

$$\sigma^2(S, t) = 2 \frac{-\frac{\partial C}{\partial t} - \mu(C - S \frac{\partial C}{\partial K}|_{K=S})}{S^2 \frac{\partial^2 C}{\partial K^2}|_{K=S}} \quad (2)$$

where C is the forward call price. This equation is of limited practical value unless some form of regularization is used. Since quoted prices are only known for a relatively small number of strikes and maturities, the presence of the curvature term in the denominator will cause the shape of the local volatility surface to be sensitive to any interpolation device used to compute the partial derivatives in the last equation. Notice that the computation of the partial derivatives in the formula for $\sigma(S, t)$ is less trivial than it seems. The forward call price depends implicitly on the Black and Scholes implied volatility surface, $\sigma_{BS}(K, T)$. This causes the choice of interpolation algorithms in both the maturity and strike directions to affect the shape of $\sigma(S, t)$.

The procedure we propose should be equivalent, in the case of call options, to the use of Dupire's equation, if great care in regularizing the price surface is taken. Unlike Dupire's formula, this procedure is general and can be applied to any type of derivative.

The approach we demonstrate is based on numerically solving the Fokker-Plank equation (FPE) for the diffusion of probability density of the underlying process. The approach is simple, reasonably fast, very easy to implement, fairly robust, and leads to very good fits with observed prices.

USING THE FOKKER-PLANK EQUATION TO COMPUTE IMPLIED VOLATILITIES

In the approach we propose, we parameterize the local volatility function in terms of a small number of parameters and solve a sequence of optimization problems to determine those parameters. Each optimization problem minimizes the difference between computed and

observed option prices of a given maturity. The optimization problems are solved sequentially in increasing maturity. There are several reasons why this approach is numerically very efficient. One reason is that solving a sequence of low-dimensional optimization problems is much easier than solving one large-dimensional optimization problem. Another reason is that by solving a sequence of optimization problems there is a great deal of reuse of computed results. Finally, the use of the FPE to price the quoted claims means that we must concern ourselves with solving very simple initial value problems, whose numerical solutions are well-behaved.

The solution approach consists in defining a functional form (in terms of the underlying price process) for the implied local volatility function, and then solving a constrained optimization problem. We *do not* parameterize the solution as a function of time. We use the following representation:

$$\sigma(S, t) = \sigma(S; c_1(t), c_2(t), \dots, c_5(t)).$$

The time dependence of the implied volatility function enters through the $c_i(t)$, which are assumed to be step-wise linear functions of time.

The simplest form of an optimization problem to imply the local volatility function would be to minimize a norm of the difference between the option prices computed with the assumed function form of the local volatility function and the observed prices. This would be an unconstrained optimization problem. This approach does not work well in practice. The reason for this is the implied volatility function is an ill-posed inverse problem, and this not only makes the solution a discontinuous function of time and difficult to work with, but also causes convergence problems as longer maturities are considered.

A way to address the issue of ill-posedness is to resort to regularization, where the smoothness of the solution is introduced explicitly into the optimization problem. Since in our case we parameterize the solution as a function of the underlying price process, the only aspect of the solution smoothness that concerns us is the partial derivative of the local volatility function with respect to time.

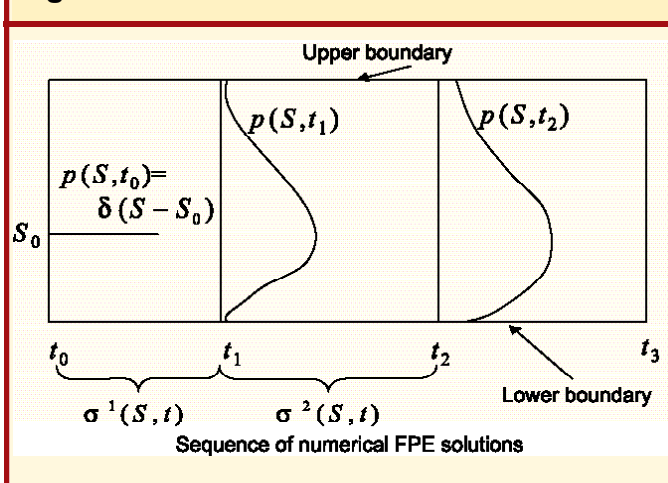
This allows us to make the following simple formulation for the optimization problem. Assume that we have observed call prices corresponding to maturities t_1, t_2, \dots, t_N . Assume that corresponding to each t_n we have $M(n)$ prices for different strikes, C_m^n . We call \hat{C}_m^n the computed value of the call with the parameters characterizing C_m^n . We define a penalty function λ , and solve the following sequence of optimization problems.

$$\min_{c_i(t_n), i=1..5} \left(\lambda \frac{\partial \sigma(c_i(t_n))}{\partial t} + \sum_{m=1}^{m=M(n)} (\hat{C}_m^n - C_m^n)^2 \right)$$

$n = 1, 2, \dots, N$

Each optimization problem uses the solution of the previous one as a starting guess. We used maturity

Figure 1



times about 0.1 yr apart. Although the number of optimization problems solved is very large (between 10 and 100), the procedure is quite effective, because each optimization problem starts out with a reasonably good guess.

The FPE describes the probability density, $p(S, t)$, of the process attaining the state S at time t , conditional on an initial state. For the price process we consider here, the FPE is given by

$$\frac{\partial p}{\partial t} + \frac{\partial (Srp)}{\partial S} - \frac{1}{2} \frac{\partial^2 (S^2 \sigma^2 p)}{\partial S^2} = 0$$

If the process S has the known value S_0 at the initial time t_0 , the initial condition of this equation is $p(S, t_0) = \delta(S - S_0)$. If the process is not known at time t_0 , but its probability density is known, then that probability density is the initial condition for the solution of the forward Fokker-Plank equation.

In our case, we will be solving this equation numerically in the various optimization problems that we discussed in the previous sections. Will consider the case of European calibration claims, but the idea can be extended easily to claims with early exercise. The important consideration is to select the most liquid claims for the particular underlying. In the first optimization problem, the initial condition will be a delta function, since the current value of the spot is known. For the subsequent optimization problems, however, the initial conditions in each each optimization problem will be the probability density at the payoff time of the previous problem. This is illustrated in Figure 1.

The boundary conditions in the absence of absorbing or reflecting barriers would have to be consistent with the behavior of probability distributions in the far field. The boundary conditions that we use must recognize that if the boundaries are located at a finite distance from the initial value of the process, probability will be transferred through the boundaries. This means that the integral of the probability density in the solution field will decrease as a function of time.

The issue of boundary conditions in a finite difference

approach is important for two reasons. One reason is that the formulation of boundary conditions affects the stability of the scheme and the accuracy of the solution, the other reason is that if not done correctly, the boundary conditions may drastically reduce the speed with which the solution is found.

There isn't a unique way of formulating boundary conditions for the numerical solution of the Fokker-Plank equation. We can consider the following possibilities.

a) Pure convection at the boundaries: In this case, we assume that we can neglect the diffusion terms at the boundaries and solve the following equation at the boundary.

$$\frac{\partial p}{\partial t} = -\frac{\partial(S\mu p)}{\partial S}$$

b) Zero curvature at the boundary: In this case we assume the following at the boundaries.

$$\frac{\partial^2 p}{\partial S^2} = 0$$

c) Full solution at the boundary: A more accurate boundary condition is to impose the full FPE at the boundary, but with all its terms discretized in terms of one-sided information inside the solution domain. This type of boundary condition has the problem that it will normally break the features required for high efficiency in implicit schemes.

We used the simple Crank-Nicholson semi-implicit discretization scheme, and found that either of the two first boundary conditions worked equally well.

To implement the optimization problems, we assume that the local volatility is given by a function of the form

$$\sigma(S, t) = C_1(t) + \frac{C_2(t)}{S^{C_5(t)}} + C_3(t)S^{C_4(t)}$$

where

$$C_i(t) = C_i(t_{n-1}) + \frac{C_i(t_n) - C_i(t_{n-1})}{t_n - t_{n-1}}(t - t_{n-1})$$

The unknown values $C_1(t_n)$, $C_2(t_n)$, $C_3(t_n)$, $C_4(t_n)$, $C_5(t_n)$ are solutions to the minimization problem. At any point in time, the first term would reflect a constant component of the local volatility. The second term would reflect the increase of volatility as the spot decreases, and the third would reflect the more gradual decrease in volatility as the spot increases.

Other functional forms are also possible, as are descriptions in terms of spline functions or Fourier decompositions.

The implementation of the successive optimization problems was done with the Broyden algorithm, a multi-dimensional Newton-type method, where the Jacobians are computed numerically. The initial guess to any given problem but the first is set to the solution of the previous problem. The initial guess for the first optimization problem is determined by trial and error.

We used interpolated quoted data for maturities between 0 and 7 years. Figure 2 shows the implied Black and Scholes volatility surface that is the basis for the computations. This data corresponds to forward options values for the Euro Stoxx 50 in mid 2000.

Figure 3 shows the implied volatility surface that corresponds to the volatility surface shown in the previous figure. These computations were done with $\lambda = 10$. The particular value of λ is not significant. The main effect of λ is to cause the surface to look more or less smooth, and to facilitate convergence. Higher values of λ will lead to smoother-looking surfaces, but the quality of the fit will not change visibly. Much lower values of λ simply cause the ripples shown in the figure to extend all over the surface.

Figure 4 shows the absolute value of the difference between the observed and the computed option values, in terms of implied Black Scholes volatility. We can see that over most of the range of interest, the agreement is within one tenth of a vol, which can be considered excellent.

CONCLUSIONS

The Fokker-Plank approach can be used in front office pricing, where accuracy and efficiency are critical. Once the implied local volatility function has been generated, which typically takes only a few minutes, the outcome can be stored and used repeatedly for fast pricing of exotic derivatives. This methodology is also useful to

Figure 2

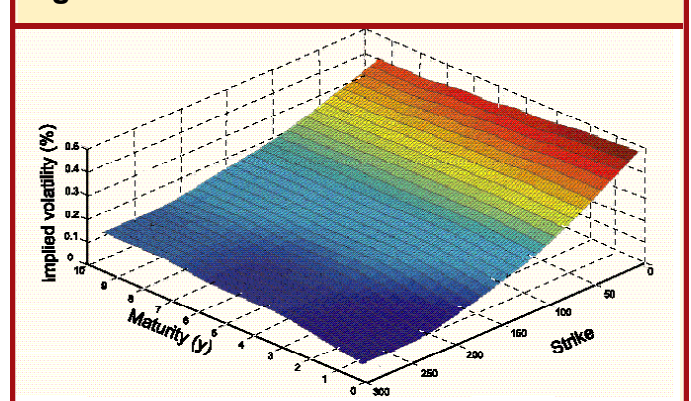

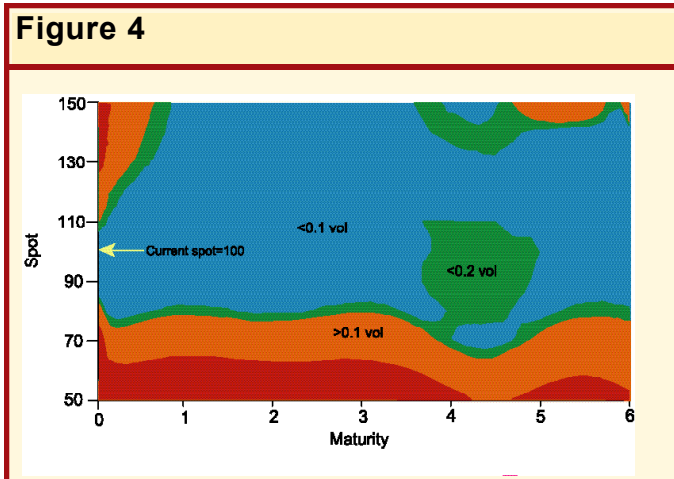


Figure 3





analyze the vega and smile risk of exotic products. Here, we report results using the Euro Stoxx 50 option market to compute the implied local volatility, but the methodology carries over easily to other underlyings such as equities and other risk classes such as interest rate derivatives and convertible bonds. In our experience, this methodology has demonstrated to be far more robust and accurate than the more conventional implied tree approaches.

Domingo Tavella is with Ocanti Associates (San Francisco, US), Wolfgang Klopfer is with Hypovereinsbank, (Munich, Germany)