Volatility Surfaces:

Theory, Rules of Thumb, and Empirical Evidence

Toby Daglish
Rotman School of Management
University of Toronto
Email: *daglish@rotman.utoronto.ca*

John Hull
Rotman School of Management
University of Toronto
Email: *hull@rotman.utoronto.ca*

Wulin Suo
School of Business
Queen's University
Email: *wsuo@business.queensu.ca*

First version: March 2001
This version: October 2002

Abstract

Implied volatilities are frequently used to quote the prices of options. The implied volatility of a European option on a particular asset as a function of strike price and time to maturity is known as the asset's volatility surface. Traders monitor movements in volatility surfaces closely. In this paper we develop a no-arbitrage condition for the evolution of a volatility surface. We examine a number of rules of thumb used by traders to manage the volatility surface and test whether they are consistent with the no-arbitrage condition and with data on the trading of options on the S&P 500 taken from the over-the-counter market.

1 Introduction

Option traders and brokers in over-the-counter markets frequently quote option prices using implied volatilities calculated from Black and Scholes (1973) and other similar models. Put–call parity implies that, in the absence of arbitrage, the implied volatility for a European call option is the same as that for a European put option when the two options have the same strike price and time to maturity. This is convenient: when quoting an implied volatility for a European option with a particular strike price and maturity date, a trader does not need to specify whether a call or a put is being considered.

The implied volatility of European options on a particular asset as a function of strike price and time to maturity is known as the volatility surface. Every day traders and brokers estimate volatility surfaces for a range of different underlying assets from the market prices of options. Some points on a volatility surface for a particular asset can be estimated directly because they correspond to actively traded options. The rest of the volatility surface is typically determined by interpolating between these points.

If the assumptions underlying Black–Scholes held for an asset, its volatility surface would be flat and unchanging. In practice the volatility surfaces for most assets are not flat and change stochastically. Consider for example equities and foreign currencies. Rubinstein (1994) and Jackwerth and Rubinstein (1996), among others, show that the implied volatilities of stock and stock index options exhibit a pronounced "skew" (that is, the implied volatility is a decreasing function of strike price). For foreign currencies this skew becomes a "smile" (that is, the implied volatility is a U-shaped function of strike price). For both types of assets, the implied volatility can be an increasing or decreasing function of the time to maturity. The volatility surface changes through time, but the general shape of the relationship between volatility and strike price tends to be preserved.

Traders use a volatility surface as a tool to value a European option when its price is not directly observable in the market. Provided there are a reasonable number of actively traded European options and these span the full range of the strike prices and times to maturity that are encountered, this approach ensures that traders price all European options consistently with the market. However, as pointed out by Hull and Suo (2002), there is no easy way to extend the approach to price path-dependent exotic options such as barrier options, compound options, and Asian options. As a result there is liable to be some model risk when these options are priced.

Traders also use the volatility surface in an *ad hoc* way for hedging. They attempt to hedge against potential changes in the volatility surface as well as against changes in the asset price. As described in Derman (1999) one popular approach to hedging against asset price movements is the "volatility-by-strike" or "sticky strike" rule. This assumes that the implied volatility for an option with a given strike price and maturity will be unaffected by changes in the underlying asset price. Another popular approach is the "volatility-by-moneyness" or "sticky delta" rule. This assumes that the volatility for a particular maturity depends only on the moneyness (that is, the ratio of the price of an asset to the strike price).

The first attempts to model the volatility surface were by Rubinstein (1994), Derman and Kani (1994), and Dupire (1994). These authors show how a one-factor model for an asset price, known as the implied volatility function (IVF) model, can be developed so that it is exactly consistent with the current volatility surface. Unfortunately, the evolution of the volatility surface under the IVF model can be unrealistic. The volatility surface given by the model at a future time is liable to be quite different from the initial volatility surface. For example, in the case of a foreign currency the initial U-shaped relationship between implied volatility and strike price is liable to evolve to one where the volatility is a monotonic increasing or decreasing function of strike price. Dumas, Fleming, and Whaley (1997) have shown that the IVF model does not capture the dynamics of market prices well. Hull and Suo (2002) have shown that it can be dangerous to use the model for the relative pricing of barrier options and plain vanilla options.

In the first part of this paper we develop a general diffusion model for the evolution of a volatility surface and derive a restriction on the specification of the model necessary for it to be a no-arbitrage model. Other researchers that have independently followed a similar approach are Ledoit and Santa Clara (1998), Schönbucher (1999), and Brace et al (2001). In addition, Britten–Jones and Neuberger (2000) produce some interesting results characterizing the set of all continuous price processes that are consistent with a given set of option prices. Our work is different from that of other researchers in that we a) investigate the implications of the no-arbitrage condition for the shapes of the volatility surfaces likely to be observed in different situations and b) examine whether the various rules of thumb that have been put forward by traders are consistent with the no-arbitrage condition. We also extend the work of Derman (1999) to examine whether the rules of thumb are supported by market data.

The rest of this paper is organized as follows. Section 2 proposes a general model for the evolution of a volatility surface and derives the no-arbitrage condition. Section 3 discusses the implications of the no-arbitrage condition. Section 4 examines a number of special cases of the model. Section 5 considers three rules of thumb used by traders and examines whether they are consistent with the no-arbitrage condition. Section 6 uses the rules of thumb to develop and test a number of hypotheses on the evolution of a volatility surface using data from the over-the-counter market. We also demonstrate how the dynamics of the volatility surface can be estimated by suing maximum likelihood estimations. Conclusions are in Section 7.

2 The Dynamics of the Implied Volatility

We suppose that the risk-neutral process followed by the price of an asset, S, is

$$\frac{dS}{S} = [r(t) - q(t)]\ dt + \sigma\ dz, \tag{1}$$

where $r(t)$ is the risk-free rate, $q(t)$ is the yield provided by the asset, σ is the asset's volatility, and z is a Wiener process. We suppose that $r(t)$ and $q(t)$ are deterministic functions of time and that σ follows a diffusion process. Our model includes the IVF model and stochastic volatility models such as Hull and White (1987), Stein and Stein (1991) and Heston (1993) as special cases.

Most stochastic volatility models specify the process for σ directly. We instead specify the processes for all implied volatilities. Define $\sigma_{TK}(t, S)$ as the implied volatility at time t of an option with strike price K and maturity T when the asset price is S and $V_{TK}(t, S)$ as the implied variance of this option $(t < T)$ so that

$$V_{TK}(t, S) = [\sigma_{TK}(t, S)]^2$$

Suppose that the process followed by V_{TK} in a risk-neutral world is

$$dV_{TK} = \alpha_{TK}\ dt + V_{TK} \sum_{i=1}^{N} \theta_{TKi}\, dz_i \tag{2}$$

where $z_1, \cdots, z_N$ are Wiener processes driving the volatility surface. Without loss of generality, we assume that these Wiener processes are uncorrelated. The z_i may be correlated with the Wiener process, z, driving the asset price in equation (1). We define ρ_i as the correlation between z and z_i. The initial volatility surface is $\sigma_{TK}(0, S_0)$ where S_0 is the initial asset price. This volatility surface can be estimated from the current ($t = 0$) prices of European call or put options and is assumed to be known.

The family of processes in equation (2) defines the multi-factor dynamics of the volatility surface. The parameter θ_{TKi} measures the sensitivity of V_{TK} to the Wiener process, z_i. In the most general form of the model the parameters α_{TK} and θ_{TKi} $(1 \leq i \leq N)$ may depend on past and present values of S, past and present values of V_{TK}, and time.

There is clearly a relationship between the instantaneous volatility $\sigma(t)$ and the volatility surface $\sigma_{TK}(t, S)$. The appendix shows that the instantaneous volatility is the limit of the implied volatility of an at-the-money option as its time to maturity approaches zero. For this purpose an at-the-money option is defined as an option where the strike price equals the forward asset price.[1] Formally:

$$\lim_{T \to t} \sigma_{TF}(t, S) = \sigma(t). \tag{3}$$

where F is the forward price of the asset at time t for a contract maturing at time T. In general the process for σ is non-Markov. This is true even when the processes in equation (2) defining the volatility surface are Markov.[2]

Define $c(S, V_{TK}, t; K, T)$ as the price of a European call option with strike price K and maturity T when the asset price, S, follows the process in equations (1) and (2). From the definition of implied volatility and the results in Black and Scholes (1973) and Merton (1973) it follows that:

$$c(S, V_{TK}, t; K, T) = e^{-\int_t^T q(\tau)\, d\tau} SN(d_1) - e^{-\int_t^T r(\tau)\, d\tau} KN(d_2)$$

where

$$d_1 = \frac{\ln(S/K) + \int_t^T [r(\tau) - q(\tau)] d\tau}{\sqrt{V_{TK}(T-t)}} + \frac{1}{2}\sqrt{V_{TK}(T-t)}$$

[1]Note that the result is not necessarily true if we define an at-the-money option as an option where the strike price equals the asset price. For example, when

$$\sigma_{TK}(t, S) = a + b\frac{1}{T-t}\ln(F/K)$$

with a and b constants, $\lim_{T \to t} \sigma_{TF}(t, S)$ is not the same as $\lim_{T \to t} \sigma_{TS}(t, S)$

[2]There is an analogy here to the Heath, Jarrow, and Morton (1992) model. When each forward rate follows a Markov process the instantaneous short rate does not in general do so.

$$d_2 = \frac{\ln(S/K) + \int_t^T [r(\tau) - q(\tau)]d\tau}{\sqrt{V_{TK}(T-t)}} - \frac{1}{2}\sqrt{V_{TK}(T-t)}$$

Using Ito's lemma equations (1) and (2) imply that the drift of c in a risk-neutral world is:

$$\frac{\partial c}{\partial t} + (r-q)S\frac{\partial c}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 c}{\partial S^2} + \alpha_{TK}\frac{\partial c}{\partial V_{TK}} + \frac{1}{2}V_{TK}^2 \frac{\partial^2 c}{\partial V_{TK}^2}\sum_{i=1}^{N}(\theta_{TKi})^2$$

$$+SV_{TK}\sigma\frac{\partial^2 c}{\partial S \partial V_{TK}}\sum_{i=1}^{N}\theta_{TKi}\rho_i$$

In the most general form of the model ρ_i, the correlation between z and z_i, is a function of past and present values of S, past and present values of V_{TK} and time. For there to be no arbitrage the process followed by c must provide an expected return of r in a risk-neutral world. It follows that

$$\frac{\partial c}{\partial t} + (r-q)S\frac{\partial c}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 c}{\partial S^2} + \alpha_{TK}\frac{\partial c}{\partial V_{TK}} + \frac{1}{2}V_{TK}^2 \frac{\partial^2 c}{\partial V_{TK}^2}\sum_{i=1}^{N}(\theta_{TKi})^2$$

$$+SV_{TK}\sigma\frac{\partial^2 c}{\partial S \partial V_{TK}}\sum_{i=1}^{N}\theta_{TKi}\rho_i = rc$$

When V_{TK} is held constant, c satisfies the Black-Scholes (1973) and Merton (1973) differential equation. As a result

$$\frac{\partial c}{\partial t} + (r-q)S\frac{\partial c}{\partial S} = rc - \frac{1}{2}V_{TK}S^2\frac{\partial^2 c}{\partial S^2}$$

It follows that

$$\frac{1}{2}S^2\frac{\partial^2 c}{\partial S^2}(\sigma^2 - V_{TK}) + \alpha_{TK}\frac{\partial c}{\partial V_{TK}} + \frac{1}{2}V_{TK}^2\frac{\partial^2 c}{\partial V_{TK}^2}\sum_{i=1}^{N}(\theta_{TKi})^2$$

$$+SV_{TK}\sigma\frac{\partial^2 c}{\partial S \partial V_{TK}}\sum_{i=1}^{N}\theta_{TKi}\rho_i = 0$$

or

$$\alpha_{TK} = -\frac{1}{2\partial c/\partial V_{TK}}\left[S^2\frac{\partial^2 c}{\partial S^2}(\sigma^2 - V_{TK}) + \frac{\partial^2 c}{\partial V_{TK}^2}V_{TK}^2\sum_{i=1}^{N}(\theta_{TKi})^2 \right.$$
$$\left. +2SV_{TK}\sigma\frac{\partial^2 c}{\partial S \partial V_{TK}}\sum_{i=1}^{N}\theta_{TKi}\rho_i\right] \quad (4)$$

The partial derivatives of c with respect to S and V_{TK} are the same as those for the Black–Scholes model:

$$\frac{\partial c}{\partial S} = e^{-\int_t^T q(\tau)\,d\tau}N(d_1)$$

$$\begin{aligned}
\frac{\partial^2 c}{\partial S^2} &= \frac{\phi(d_1)e^{-\int_t^T q(\tau)\,d\tau}}{S\sqrt{V_{TK}(T-t)}} \\
\frac{\partial c}{\partial V_{TK}} &= \frac{Se^{-\int_t^T q(\tau)\,d\tau}\phi(d_1)\sqrt{T-t}}{2\sqrt{V_{TK}}} \\
\frac{\partial^2 c}{\partial V_{TK}^2} &= \frac{Se^{-\int_t^T q(\tau)\,d\tau}\phi(d_1)\sqrt{T-t}}{4V_{TK}^{3/2}}(d_1d_2-1) \\
\frac{\partial^2 c}{\partial S\partial V_{TK}} &= -\frac{e^{-\int_t^T q(\tau)\,d\tau}\phi(d_1)d_2}{2V_{TK}}
\end{aligned}$$

where ϕ is the density function of the standard normal distribution:

$$\phi(x) = \frac{1}{2\pi}\exp\left(-\frac{x^2}{2}\right), \quad -\infty < x < \infty.$$

Substituting these relationships into equation (4) and simplifying we obtain

$$\alpha_{TK} = \frac{1}{T-t}(V_{TK}-\sigma^2) - \frac{V_{TK}(d_1d_2-1)}{4}\sum_{i=1}^{N}(\theta_{TKi})^2 + \sigma d_2\sqrt{\frac{V_{TK}}{T-t}}\sum_{i=1}^{N}\theta_{TKi}\rho_i \quad (5)$$

Equation (5) provides an expression for the risk-neutral drift of an implied variance in terms of its volatility. The first term on the right hand side is the drift arising from the difference between the implied variance and the instantaneous variance. The second term arises from the part of the uncertainty about future volatility that is uncorrelated with the asset price. The third term arises from the correlation between the asset price and its volatility.

The first term can be understood by considering the situation where the instantaneous variance, σ^2, is a deterministic function of time. The variable V_{TK} is then also a function of time and

$$V_{TK} = \frac{1}{T-t}\int_t^T \sigma(\tau)^2 d\tau$$

Differentiating with respect to time we get

$$\frac{dV_{TK}}{dt} = \frac{1}{T-t}[V_{TK}-\sigma(t)^2]$$

This is the first term.

The analysis can be simplified slightly by considering the variable $\hat{V}_{TK}$ instead of V where $\hat{V}_{TK} = (T-t)V_{TK}$. Because

$$d\hat{V}_{TK} = -V_{TK}dt + (T-t)dV_{TK}$$

it follows that

$$d\hat{V}_{TK} = \left[-\sigma^2 - \frac{\hat{V}_{TK}(d_1d_2-1)}{4}\sum_{i=1}^{N}(\theta_{TKi})^2 + \sigma d_2\sqrt{\hat{V}_{TK}}\sum_{i=1}^{N}\theta_{TKi}\rho_i\right] dt$$
$$+\hat{V}_{TK}\sum_{i=1}^{N}\theta_{TKi}dz_i$$

3 Implications of the No-Arbitrage Condition

Equation (5) provides a no-arbitrage condition for the drift of the implied variance as a function of its volatility. In this section we examine the implications of this no-arbitrage condition. In the general case where V_{TK} is nondeterministic the first term in equation (5) is mean fleeing; that is it provides negative mean reversion. This negative mean reversion becomes more pronounced as the option approaches maturity. For a viable model the θ_{TKi} must be complex functions that in some way offset this negative mean reversion. Determining the nature of these functions is not easy. However, it is possible to make some general statements about the volatility smiles that are consistent with stable models.

3.1 The Zero Correlation Case

Consider first the case where all the ρ_i are zero so that the third term in equation (5) disappears. As before we define F as the forward value of the asset for a contract maturing at time T so that

$$F = Se^{\int_t^T [r(\tau)-q(\tau)]d\tau}$$

¿From equation (5) the drift of $V_{TK} - V_{TF}$ is

$$\frac{1}{T-t}(V_{TK}-V_{TF}) - \frac{V_{TK}(d_1d_2-1)}{4}\sum_{i=1}^{N}(\theta_{TKi})^2 - \frac{V_{TF}[1+V_{TF}(T-t)/4]}{4}\sum_{i=1}^{N}(\theta_{TFi})^2$$

Suppose that $d_1d_2 - 1 > 0$ and $V_{TK} < V_{TF}$. In this case each term in the drift of $V_{TK} - V_{TF}$ is negative. As a result $V_{TK} - V_{TF}$ tends to get progressively more negative and the model is unstable with negative values of V_{TK} being possible. We deduce from this that V_{TK} must be greater than V_{TF} when $d_1d_2 > 1$. The condition $d_1d_2 > 1$ is satisfied for very large and very small values of K. It follows that the case where the ρ_i are zero can be consistent with the U-shaped volatility smile. It cannot be consistent with a upward or downward sloping smile because in these cases $V_{TK} < V_{TF}$ for either very high or very low values of K.

Our finding is consistent with a result in Hull and White (1987). These authors show that when the instantaneous volatility is independent of the asset price, the price of a European option is the Black–Scholes price integrated over the distribution of the average variance. They demonstrate that when $d_1 d_2 > 1$ a stochastic volatility tends to increase an option's price.

A U-shaped volatility smile is commonly observed for options on a foreign currency. Our analysis shows that this is consistent with the empirical result that the correlation between implied volatilities and the exchange rate is close to zero (see, for example, Bates (1996)).

3.2 Volatility Skews

Consider next the situation where the volatility is a declining function of the strike price. The variance rate V_{TK} is greater than V_{TF} when K is very small and less than V_{TF} when K is very large. When we do not make the zero-correlation assumption the drift of $V_{TK} - V_{TF}$ is

$$\frac{1}{T-t}(V_{TK} - V_{TF}) - \frac{V_{TK}(d_1 d_2 - 1)}{4}\sum_{i=1}^{N}(\theta_{TKi})^2 - \frac{V_{TF}(1 + V_{TF}(T-t)/4)}{4}\sum_{i=1}^{N}(\theta_{TFi})^2$$
$$+\sigma d_2\sqrt{\frac{V_{TK}}{T-t}}\sum_{i=1}^{N}\theta_{TKi}\rho_i + \sigma\frac{V_{TF}}{2}\sum_{i=1}^{N}\theta_{TFi}\rho_i$$

When $K > F$, $V_{TK} - V_{TF}$ is negative and the effect of the first three terms is to provide a negative drift as before. For a stable model we require the last two terms to give a positive drift. As K increases and we approach option maturity the first of the last two terms dominates the second. Because $d_2 < 0$ we must have

$$\sum_{i=1}^{N}\theta_{TKi}\rho_i < 0 \tag{6}$$

when K is very large. The instantaneous covariance of the asset price and its variance is

$$\sigma\sum_{i=1}^{N}\theta_{TKi}\rho_i$$

Because $\sigma > 0$ it follows that when K is large the asset price must be negatively correlated with its variance.

Equities provide an example of a situation where there is a volatility skew of the sort we are considering. As has been well documented by authors such as Christie (1982), the volatility of an equity price tends to be negatively correlated with the equity price. This is consistent with the result we have just presented.

3.3 Other Results

Consider the situation where volatility is an increasing function of the strike price. (This is the case for options on some commodity futures.) A similar argument to that just given shows that the no-arbitrage relationship implies that the volatility of the variable should be positively correlated with the level of the variable.

Another possibility for the volatility smile is an inverted-U-shaped pattern. In this case a similar analysis to that given above shows that for a stable model

$$\sum_{i=1}^{N} \theta_{TKi}\rho_i$$

must be less than zero when K is large and greater than zero when K is small. It is difficult to see how this can be so without the stochastic terms in the processes for the V_{TK} having a form that quickly destroys the inverted-U-shaped pattern.

4 Special Cases

In this section we consider a number of special cases of the model developed in Section 2.

Case 1: V_{TK} is a deterministic function only of t, T, and K

In this situation $\theta_{TKi} = 0$ for all $i \geq 1$ so that from equation (2)

$$dV_{TK} = \alpha_{TK}\ dt,$$

Also from equation (5)

$$\alpha_{TK} = \frac{1}{T-t}\left[V_{TK} - \sigma^2\right]$$

so that

$$dV_{TK} = \frac{1}{T-t}\left[V_{TK} - \sigma^2\right]\ dt$$

which can be written as

$$(T-t)dV_{TK} - V_{TK}dt = -\sigma^2\ dt$$

or

$$\sigma^2 = -\frac{d\left[(T-t)V_{TK}\right]}{dt} \tag{7}$$

This shows that σ is a deterministic function of time. The only model that is consistent with V_{TK} being a function only of t, T, and K is therefore the model where the instantaneous volatility, σ, is a function only of time. This is Merton's (1973) model.

In the particular case where V_{TK} depends only on T and K, equation (7) shows that $V_{TK} = \sigma^2$ and we get the Black-Scholes constant-volatility model.

Case 2: V_{TK} is independent of the asset price, S .

In this situation ρ_i is zero and equation (5) becomes

$$\alpha_{TK} = \frac{1}{T-t}(V_{TK} - \sigma^2) - \frac{V_{TK}(d_1 d_2 - 1)}{4} \sum_{i=1}^{N} (\theta_{TKi})^2$$

Both α_{TK} and the θ_{TKi} must be independent of S. Because d_1 and d_2 depend on S we must have θ_{TKi} equal zero for all i. Case 2 therefore reduces to Case 1. The only model that is consistent with V_{TK} being independent of S is therefore Merton's (1973) model where the instantaneous volatility is a function only of time.

Case 3: V_{TK} is a deterministic function of t, T, and K/S, or equivalently, K/F.

In this situation

$$V_{TK} = G\left(T, t, \frac{K}{F}\right),$$

where G is a deterministic function and as before F is the forward price of S. From equation (3) the spot instantaneous volatility, σ, is given by

$$\sigma^2 = G\left(t, t, 1\right),$$

This is a deterministic function of time. It follows that, yet again, the model reduces to Merton's (1973) deterministic volatility model.

Case 4: V_{TK} is a deterministic function of t, T, S and K.

In this situation we can write

$$V_{TK} = G(T, t, F, K), \tag{8}$$

where G is a deterministic function. From equation (3), the instantaneous volatility, σ is given by

$$\sigma^2 = \lim_{T \to t} G(T, t, F, F)$$

This shows that the instantaneous volatility is a deterministic function of F and t. Equivalently it is a deterministic function of underlying asset price, S, and t. It follows that the model reduces to the IVF model. Writing σ as $\sigma(S,t)$, Dupire (1994) and Andersen and Brotherton–Ratcliffe (1998) show that

$$[\sigma(K,T)]^2 = 2\frac{\partial c/\partial T + q(T)c + K[r(T) - q(T)]\partial c/\partial K}{K^2(\partial^2 c/\partial K^2)}$$

where c is here regarded as a function of S, K and T for the purposes of taking partial derivatives.

5 Rules of Thumb

A number of rules of thumb have been proposed about volatility surfaces. These rules of thumb fall into two categories. In the first category are rules concerned with the way in which the volatility surface changes through time. They are useful in the calculation of the Greek letters such as delta and gamma. In the second category are rules concerned with the relationship between the volatility smiles for different option maturities at a point in time. They are useful in creating a complete volatility surface when market prices are available for a relatively small number of options. In this section we consider three different rules of thumb: the sticky strike rule, the sticky delta rule and the square root of time rule. The first two of these rules are in the first category and provide a basis for calculating Greek letters. The third rule is in the second category and assists with the mechanics of "filling in the blanks" when a complete volatility surface is being produced.

5.1 The Sticky Strike Rule

The sticky strike rule assumes that the implied volatility of an option is independent of the asset price. This is an appealing assumption because it implies that the sensitivity of the price of an option to S is

$$\frac{\partial c}{\partial S}$$

where for the purposes of calculating the partial derivative the option price, c, is considered to be a function of S, V_{TK}, and t. The assumption enables the Black–Scholes formulas to be used to calculate delta with the volatility parameter set equal to the option's implied volatility. The same is true of gamma.

Under the most basic form of the sticky strike rule the implied volatility of an option is assumed to remain the same for the whole of its life. The variance V_{TK} is a function only of K and T. The analysis in Case 1 of the previous section shows that the only version of this model that is internally consistent is the model where the volatilities of all options are the same and constant. This is the original Black and Scholes (1973) model.

A rather more sophisticated version of the sticky strike rule is where V_{TK} is independent of S, but possibly dependent on other stochastic variables. As shown in Case 2 of the previous section, the only version of this model that is internally consistent is the model where the instantaneous volatility of the asset price is a function only of time. This is

Merton's (1973) model.

When the instantaneous volatility of the asset price is a function only of time, all European options with the same maturity have the same implied volatility. We conclude that all versions of the sticky strike rule are inconsistent with any type of volatility smile or volatility skew. If a trader prices options using different implied volatilities and the volatilities are independent of the asset price, there must be arbitrage opportunities.

5.2 The Sticky Delta Rule

An alternative to the sticky strike rule is the sticky delta rule. This assumes that the implied volatility of an option depends on S and K through its dependence on the moneyness variable, K/S. The delta of a European option in a stochastic volatility model is

$$\Delta = \frac{\partial c}{\partial S} + \frac{\partial c}{\partial V_{TK}}\frac{\partial V_{TK}}{\partial S}$$

Again, for the purposes of calculating partial derivatives the option price c is considered to be a function of S, V_{TK}, and t. The first term in this expression is the delta calculated using Black-Scholes with the volatility parameter set equal to implied volatility. In the second term, $\partial c/\partial V_{TK}$ is positive. It follows that, if V_{TK} is a declining (increasing) function of the strike price, it is an increasing (declining) function of S and Δ is greater than (less than) that given by Black-Scholes. For equities V_{TK} is a declining function of K and so the Black-Scholes delta understates the true delta. For an asset with a U-shaped volatility smile Black-Scholes understates delta for low strike prices and overstates it for high strike prices.

In the most basic form of the sticky delta rule the implied volatility is assumed to be a deterministic function of K/S and $T - t$. Case 3 in the previous section shows that the only version of this model that is internally consistent is Merton's (1973) model where the instantaneous volatility of the asset price is a function only of time. Again we find that the model is inconsistent with any type of volatility smile or volatility skew. For no arbitrage, implied volatilities must be independent of S and K.

A more general version of the sticky delta rule is where the process for V_{TK} depends on K, S, T, and t only through its dependence on K/S and $T - t$. We will refer to this as the *generalized sticky delta model.* Models of this type can be consistent with the no-arbitrage condition. This is because equation (5) shows that if each θ_{TKi} depends on K, S, T, and t

only through a dependence on K/S and $T-t$, the same is true of α_{TK}.

Many traders argue that a better measure of moneyness than K/S is K/F where as before F is the forward value of S for a contract maturing at time T. A version on the sticky delta rule often used by traders is that $\sigma_{TK}(S,t) - \sigma_{TF}(S,t)$ is a function only of K/F and $T-t$. Here it is the excess of the volatility over the at-the-money volatility, rather than the volatility itself, which is assumed to be a deterministic function of the moneyness variable, K/F. This form of the sticky delta rule allows the overall level of volatility to change through time and the shape of the volatility term structure to change, but when measured relative to the at-the-money volatility, the volatility is dependent only on K/S and $T-t$. We will refer to this model as the *relative sticky delta model.* If the at-the-money volatility is stochastic, but independent of S, the model is a particular case of the generalized sticky delta model just considered.

5.3 The Square Root of Time Rule

A rule that is sometimes used by traders is what we will refer to as the "square root of time rule". This is described in Natenberg (1994) and Hull (2002). It provides a specific relationship between the volatilities of options with different strike prices and times to maturity at a particular time. One version of rule is

$$\frac{\sigma_{TK}(S,t)}{\sigma_{TF}(S,t)} = \Phi\left(\frac{\ln(K/F)}{\sqrt{T-t}}\right)$$

where Φ is a function, and F is the forward price at time t of the underlying with a contract maturity of T. An alternative that we will use is

$$\sigma_{TK}(S,t) - \sigma_{TF}(S,t) = \Phi\left(\frac{\ln(K/F)}{\sqrt{T-t}}\right) \tag{9}$$

We will refer version of the rule where Φ does not change through time as the *stationary square root of time model* and the version of the rule where the form of the function Φ changes stochastically as the *stochastic square root of time model.*

The square root of time model (whether stationary or stochastic) simplifies the specification of the volatility surface. If we know

1. The volatility smile for options that mature at one particular time T^*, and
2. At-the-money volatilities for other maturities

we can compute the complete volatility surface. Suppose that F^* is the forward price of the asset for a contract maturing at T^*. We can compute the volatility smile at time T from that at time T^* using the result that

$$\sigma_{TK}(S,t) - \sigma_{TF}(S,t) = \sigma_{T^*K^*}(S,t) - \sigma_{T^*F^*}(S,t)$$

where

$$\frac{\ln(K/F)}{\sqrt{T-t}} = \frac{\ln(K^*/F^*)}{\sqrt{T^*-t}}$$

or

$$K^* = F^* \left(\frac{K}{F}\right)^{\sqrt{(T^*-t)/(T-t)}}$$

If the at-the-money volatility is assumed to be stochastic, but independent of S, the stationary square root of time model is a particular case of the relative sticky strike model and the stochastic square root of time model is a particular case of the generalized sticky strike model.

6 Empirical Tests

As pointed out by Derman (1999) apocryphal rules of thumb for describing how volatility smiles and skews change may not be confirmed by data. Derman's research looks at options on the S&P 500 during the period September 1997 to October 1998 and considers the sticky-strike and sticky delta rules as well as a more complicated rule based on the IVF model. He finds subperiods during which each of rules appears to explain the data best.

Derman's results are based on options with three-month maturities. Options on the S&P 500 regularly trade with maturities up to five years in the over-the-counter market. In this section we consider a wide range of option maturities to try and obtain a more complete picture of the rules governing the behavior of the volatility surface. We also investigate the factors deriving the volatility surface. The multiple option maturities also allow us to test the square root of time rule. To the best of our knowledge there are no previous tests of this rule in the literature.

6.1 Data

The data we use are monthly volatility surfaces for 47 months (June 1998 to April 2002). The data for June 1998 is shown in Table 1. Six maturities are considered ranging from six months to five years. Seven values of K/S are considered ranging from 80 to 120. A total of 42 points on the volatility surface are therefore provided each month and the total number of volatilities in our data set is $42 \times 47 = 1974$. As illustrated in Table 1, implied volatilities for the S&P 500 exhibit a volatility skew with $\sigma_{TK}(t, S)$ being a decreasing function of K.

[Table 1 about here.]

The data was supplied to us by Totem Market Valuations Limited with the kind permission of a selection of Totem's major bank clients. Totem collects implied volatility data in the form shown in Table 1 from a large number of dealers each month. These dealers are market makers in the over-the-counter market. Totem uses the data in conjunction with appropriate averaging procedures to produce an estimate of the mid-market implied volatility for each cell of the table and returns these estimates to the dealers. This enables dealers to check whether their valuations are in line with the market. Our data consists of the estimated mid-market volatilities returned to dealers. The data produced by Totem are considered by market participants to be more accurate than either the volatility surfaces produced by brokers or those produced by any one individual bank.

6.2 Estimation of the Volatility Factors

We have assumed that the volatility surface is driven a certain number of factors. In this part, we illustrate how how to use maximum likelihood to estimate these factors, and the number of factors needed to describe the volatility surface. For this purpose, we consider the special case where $\theta_{TK}(t, S)$ is of such a form that we may consider $\theta_{(T-t),F/K}$ to be a constant, consistent with the sticky delta rule of Section 5.1. We use our data to generate volatility matrices where the data is spaced according to F/K rather than S/K, as originally reported. We create this data, and the attendant changes in volatility, using interpolation. Given this, we can consider a discretized process for $\log(V_{(T-t),F/K})$:

$$\Delta \log(V_{(T-t),F/K,t}) = \hat{\alpha}_{(T-t),F/K}\Delta t + \epsilon_{(T-t),F/K,t}$$

where

$$\hat{\alpha}_{(T-t),F/K} = \frac{1}{T-t}(1 - \frac{\sigma^2}{V_{(T-t),F/K}})$$

$$-\frac{(d_1 d_2 + 1)}{4} \sum_{i=1}^{N} \left[\theta^2_{(T-t),F/K,i} + \sum_{j \neq i} \rho_i \rho_j \theta_{(T-t),F/K,i} \theta_{(T-t),F/K,j} \right]$$
$$+\sigma d_2 \frac{1}{\sqrt{V_{(T-t),F/K}(T-t)}} \sum_{i=1}^{N} \theta_{(T-t),F/K,i} \rho_i$$

and the errors are distributed such that

$$E(\epsilon_{(T-t),F/K,t}) = 0$$

and

$$E(\epsilon_{(T-t)_i,(F/K)_i,t} \epsilon_{(T-t)_j,(F/K)_j,t}) = \sqrt{\Delta t} \sum_{k=1}^{N} \left[\theta^2_{(T-t)_i,(F/K)_i,k} + \sum_{l \neq k} \theta_{(T-t)_i,(F/K)_i,k} \theta_{(T-t)_j,(F/K)_j,l} \rho_k \rho_l \right] + \begin{cases} \sigma^2_\epsilon & \text{if } i = j. \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

Note that this gives our data a block diagonal variance-covariance matrix. Changes in implied volatilities at the same time are correlated with each other, but no two volatility movements at different times are correlated. Assuming that the ϵ's are normally distributed, we can find the likelihood function for each volatility matrix. For $\theta_{(T-t),(F/K),i}$ and ρ_i $i = 1 \ldots N$, along with the residual standard deviation (σ^2_ϵ) :

$$L = \sum_{i=1}^{M} -\epsilon'_{t_i} V_{t_i} \epsilon_{t_i} - \log \det(V_{t_i})$$

where V is the variance covariance matrix described by (10) and $t_1 \ldots t_M$ are the times on which volatility surfaces are observed. The ability to find the likelihood for each volatility matrix separately follows from the block-diagonality of the variance covariance matrix. Estimation consists of finding $\theta_{(T-t),(F/K),i}$ and ρ_i $i = 1 \ldots N$, and the residual standard deviation (σ^2_ϵ) so as to minimize L.

Fitting a three factor model ($N = 4$) we find the following factors:

[Figure 1 about here.]

The first factor has $\rho_1 = 1.0000$, the second factor has $\rho_2 = 0.965$, while the third factor has $\rho_3 = 0.369$, and the fourth factor has $\rho_4 = 0.946$.

[Table 2 about here.]

[Table 3 about here.]

[Table 4 about here.]

[Table 5 about here.]

From these results, we can see that the first factor is common for both the underlying index level and its implied volatility surface. The uniformly positive signs of the coefficients suggest that the factor corresponds to a parallel shift to the volatility surface. For the second factor, the signs for those long term, out-of-the-money put options are negative, while they are positive for the in-the-money put options with short maturity. This factor can be regarded as an effect on the slope of the volatility surface. The magnitude of the coefficients are also smaller than those for the first factor. Similar observations can be made for the third and fourth factors.

With four factors, the standard deviation of remaining errors is 0.0153.

6.3 Tests of the Volatility Rules

We first used the data to test the sticky strike rule. The most basic version of the sticky strike rule, where the implied volatility is a function only of K and T, may be plausible in the exchange-traded market where the exchange defines a handful of options that trade and traders anchor on the volatility they first use for any one of these options. Our data comes from the over-the-counter market. It is difficult to see how the basic version of the sticky strike rule can apply in that market because there is continual trading in options with many different strike prices and times to maturity. A more plausible version of the rule in the over-the-counter market is that the implied volatility is a function only of K and $T - t$.

We tested this version of the sticky strike rule using

$$\sigma_{TK}(t, S) = a_0 + a_1 K + a_2 K^2 + a_3(T - t) + a_4(T - t)^2 + a_5 K(T - t) + \epsilon \qquad (11)$$

where the a_i are constant and ϵ is a normally distributed error term. The terms on the right hand side of this equation can be thought of as the first few terms in a Taylor series expansion of a general function of K and $T-t$. As shown in Table 2, the model is supported by the data, but has an R^2 of only 27%.

[Table 6 about here.]

We now move on to test the sticky delta rule. The version of the rule we consider is the relative sticky strike model where $\sigma_{TK}(t,S) - \sigma_{TF}(t,S)$ is a function of K/F and $T-t$. The model we test is

$$\begin{aligned}\sigma_{TK}(t,S) - \sigma_{TF}(t,S) &= b_0 + b_1 \ln\left(\frac{K}{F}\right) + b_2\left[\ln\left(\frac{K}{F}\right)\right]^2 + b_3(T-t) \\ &\quad + b_4(T-t)^2 + b_5 \ln\left(\frac{K}{F}\right)(T-t) + \epsilon \end{aligned} \tag{12}$$

where the b_i are constants and ϵ is a normally distributed error term. The model is analogous to the one used to test the sticky strike rule. The terms on the right hand side of this equation can be thought of as the first few terms in a Taylor series expansion of a general function of $\ln(K/F)$ and $T-t$. The results are shown in Table 7. In this case the R^2 is much higher at 94.93%.

[Table 7 about here.]

If two models have equal explanatory power, then the observed ratio of the two models' squared errors should be distributed $F(N_1, N_2)$ where N_1 and N_2 are the number of degrees of freedom in the two models. When comparing the sticky strike and relative sticky delta models using a two tailed test, this statistic must be greater than 1.12 or less than 0.89 for significance at the 1% level. The value of the statistic is 32.6 indicating that we can overwhelmingly reject the hypothesis that the models have equal explanatory power. The relative sticky delta model in equation (12) can explain the volatility surfaces in our data much better than the sticky strike model in equation (11).

The third model we test is the version of the stationary square root of time rule where the function Φ in equation (9) does not change through time so that $\sigma_{TK}(t,S) - \sigma_{TF}(t,S)$ is a known function of $\ln(K/S)/\sqrt{T-t}$. Using a similar Taylor Series expansion to the other models we test

$$\sigma_{TK}(t,S) - \sigma_{TF}(t,S) = c_1 \frac{\ln(K/F)}{\sqrt{T-t}} + c_2 \frac{[\ln(K/F)]^2}{T-t} + \epsilon \tag{13}$$

where c_1 and c_2 are constants and ϵ is a normally distributed error term. The results for this model are shown in Table 8. In this case the R^2 is 97.12%. This is somewhat better than the R^2 for model in (12) even though the model in equation (13) involves two parameters and the the one in equation (12) involves six parameters.

[Table 8 about here.]

We can calculate a ratio of sums of squared errors to compare the stationary square root of time model in equation (13) to the relative sticky delta model in equation (12). In this case, the ratio is 1.11. When a two tailed test is used it is not quite possible to reject the hypothesis that the two models have equal explanatory power at the 1% level, but it is possible to reject this hypothesis at the 5% level. We conclude that equation (13) is an improvement over equation (12).

In the stochastic square root of time rule, the functional relationship between $\sigma_{TK}(t,S)-\sigma_{TF}(t,S)$ and

$$\frac{\ln(K/F)}{\sqrt{T-t}}$$

changes stochastically through time. A model capturing this is

$$\sigma_{TK}(t,S)-\sigma_{TF}(t,S)=c_1(t)\frac{\ln(K/F)}{\sqrt{T-t}}+c_2(t)\frac{[\ln(K/F)]^2}{T-t}+\epsilon \tag{14}$$

where $c_1(t)$ and $c_2(t)$ are stochastic. To provide a test of this model we fitted the model in equation (13) to the data on a month by month basis. When the model in equation (14) is compared to the model in equation (13) the ratio of sums of squared errors statistic is 2.01 indicating that the stochastic square root of time model does provide a significantly better fit to the data that the stationary square root of time model at the 1% level.

[Table 9 about here.]

The coefficient, c_1, in the monthly tests of the square root of time rule is always significantly different from zero with a very high level of confidence. Interestingly, in 20 of the 47 months it was not possible to reject the hypothesis that $c_2 = 0$ at the 5% level suggesting an approximately linear relationship between $\sigma_{TK}(t,S)-\sigma_{TF}(t,S)$ and[3]

$$\frac{\ln(K/F)}{\sqrt{T-t}}$$

Figure 1 shows the level of c_1 and the S&P 500 for the period covered by our data. The coefficient of correlation between changes in c_1 and changes in the level of the S&P 500 is 21.5 percent. This is statistically significantly different from zero at the 10% confidence level. An increase (reduction) in c_1 corresponds to a decrease (increase) in the skew. The positive correlation can be viewed as an extension of the crashophobia phenomenon identified by

[3]The approximate linearity of the volatility skew for S&P 500 options has been mentioned by a number of researchers including Derman (1999).

Rubinstein (1994). When the level of the S&P 500 decreases (increases) investors become more concerned about the possibility of a crash and the volatility skew becomes more pronounced (less pronounced).

[Figure 2 about here.]

7 Summary

It is a common practice in the over-the-counter markets to quote option prices using their Black–Scholes implied volatilities. In this paper we have developed a model of the evolution of implied volatilities and produced a no-arbitrage condition that must be satisfied by the volatilities. Our model is exactly consistent with the initial volatility surface, but more general than the IVF model of Rubinstein (1994), Derman and Kani (1994), and Dupire (1994). The no-arbitrage condition leads to the conclusions that a) when the volatility is independent of the asset price there must be a U-shaped volatility smile and b) when the implied volatility is a decreasing (increasing) function of the asset price there must be a negative (positive) correlation between the volatility and the asset price.

A number of rules of thumb have been proposed for how traders manage the volatility surface. These are the sticky strike, sticky delta, and square root of time rules. Some versions of these rules are clearly inconsistent with the no-arbitrage condition; for other versions of the rules the no-arbitrage condition can in principle be satisfied.

Our empirical tests of the rules of thumb using 47 months of volatility surfaces for the S&P 500 show that the relative sticky delta model (where the excess of the implied volatility of an option over the corresponding at-the-money volatility is a function of moneyness) outperforms the sticky strike rule. Also, the stochastic square root of time model outperforms the relative sticky delta rule.

A Proof of Equation (3)

For simplicity of notation, we assume that r and q are constants. Define $F(\tau)$ as the forward price at time τ for a contract maturing at time $t + \Delta t$ so that

$$F(t) = S(t)e^{(r-q)\Delta t}$$

The price at time t of a call option with strike price $F(t)$ and maturity $t+\Delta t$ is given by

$$c(S(t), V_{t+\Delta t, F(t)}, t; F(t), t+\Delta t) = e^{-q\Delta t} S(t)[N(d_1) - N(d_2)]$$

where in this case

$$d_1 = -d_2 = \frac{\sqrt{\Delta t}}{2} \sigma_{t+\Delta t, F(t)}(t, S(t))$$

The call price can be written

$$c(S(t), V_{t+\Delta t, F(t)}, t; F(t), t+\Delta t) = e^{-q\Delta t} S(t) \int_{-d_1}^{d_1} \phi(x)\ dx$$

where

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right)$$

For some $\bar{x}$

$$\int_{-d_1}^{d_1} \phi(x)\ dx = 2d_1 \phi(\bar{x})$$

and the call price is therefore given by

$$c(S(t), V_{t+\Delta t, F(t)}, t; F(t), t+\Delta t) = 2e^{-q\Delta t} S(t) d_1 \phi(\bar{x})$$

or

$$c(S, V_{t+\Delta t, F(t)}, t; F(t), t+\Delta t) = e^{-q\Delta t} S(t) \phi(\bar{x}) \sigma_{t+\Delta t, F(t)}(t, S(t)) \sqrt{\Delta t} \tag{15}$$

The process followed by S is

$$\frac{dS}{S} = (r-q)\, dt + \sigma\, dz$$

Using Ito's lemma

$$dF = \sigma F\, dz$$

When terms of order higher that $(\Delta t)^{1/2}$ are ignored

$$F(t+\Delta t) - F(t) = \sigma(t) F(t)[z(t+\Delta t) - z(t)]$$

It follows that

$$\lim_{\Delta t \to 0} \frac{1}{\sqrt{\Delta t}} E[F(t+\Delta t) - F(t)]^+ = \lim_{\Delta t \to 0} \frac{1}{\sqrt{\Delta t}} \sigma(t) F(t) E[z(t+\Delta t) - z(t)]^+$$

where E denotes expectations under the risk-neutral measure.

Because

$$E[F(t+\Delta t) - F(t)]^+ = e^{r\Delta t} c(S(t), V_{t+\Delta t, F(t)}, t; F(t), t+\Delta t)$$

and

$$\lim_{\Delta t \to 0} \frac{1}{\sqrt{\Delta t}} E[z(t+\delta t) - z(t)]^+ = \frac{1}{\sqrt{2\pi}}$$

it follows that

$$\lim_{\Delta t \to 0} \frac{1}{\sqrt{\Delta t}} e^{r\Delta t} c(S(t), V_{t+\Delta t, F(t)}, t; F(t), t+\Delta t) = \frac{\sigma(t)F(t)}{\sqrt{2\pi}}$$

Substituting from equation (15)

$$\lim_{\Delta t \to 0} e^{r\Delta t} e^{-q\Delta t} S(t)\phi(\bar{x})\sigma_{t+\Delta t, F(t)}(t, S(t)) = \frac{\sigma(t)F(t)}{\sqrt{2\pi}}$$

As Δt tends to zero, $\phi(\bar{x})$ tends to $1/\sqrt{2\pi}$ so that

$$\lim_{\Delta t \to 0} \sigma_{t+\Delta t, F(t)}(t, S(t)) = \sigma(t)$$

This is the required result.

References

Black, F. and M.S. Scholes (1973), "The pricing of options and corporate liabilities," *Journal of Political Economy*, 81, pp637–659

Bates, D. S. (1996), "Jumps and Stochastic Volatility: Exchange Rate Process Implicit in Deutsche Mark Options," *Review of Financial Studies*, 9, 1, pp69–107

Brace, A., B. Goldys, F. Klebaner, and R. Womersley (2001), "Market model of stochastic volatility with applications to the BGM model," Working paper S01-1, Dept of Statistics, University of New South Wales

Britten–Jones, M., and A. Neuberger (2000) "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance*, 55, 2 (2000), pp839–866

Christie, A. A. (1982), "The stochastic behavior of common stock variances: Value, leverage, and interest rate effects," *Journal of Financial Economics*, 10, 4, pp407–432

Derman, E. (1999), "Regimes of volatility," *Quantitative Strategies Research Notes*, Goldman Sachs, New York, NY. Also in *Risk Magazine*, April, 1999

Derman, E. and I. Kani (1994a), "The volatility smile and its implied tree," *Quantitative Strategies Research Notes*, Goldman Sachs, New York, NY

Dumas, B., J. Fleming and R.E. Whaley (1997), "Implied volatility functions: empirical tests," *Journal of Finance* 6, pp2059–2106

Dupire, B. (1994), "Pricing with a smile," *Risk*, 7, pp18–20

Heath, D., R. Jarrow, and A. Morton (1992), "Bond Pricing and the Term Structure of the Interest Rates: A New Methodology," *Econometrica*, 60, 1, 77–105

Heston, S.L. (1993), "A closed-form solution for options with stochastic volatility applications to bond and currency options," *Review of Financial Studies*, 6, pp327–343

Hull, J. (2002) *Options, Futures, and Other Derivatives*, 5th Edition, Prentice Hall, Englewood Cliffs, NJ

Hull, J. and W. Suo (2002), "A methodology for assessing model risk and its application to the implied volatility function model," *Journal of Financial and Quantitative Analysis*, 37, 2, pp297–318

Hull, J. and A. White (1987), "The pricing of options with stochastic volatility," *Journal of Finance*, 42, pp281–300

Jackwerth, J.C. and M. Rubinstein (1996), "Recovering probabilities from option prices", *Journal of Finance*, 51, pp1611–1631

LeDoit O. and P. Santa Clara (1998) "Relative pricing of options with stochastic volatility" Working Paper, Anderson Graduate School of Management, University of California, Los Angeles

Merton, R. C. (1973), "Theory of rational option pricing," *Bell Journal of Economics and Management Science*, 4, pp141–183

Natenberg S. (1994) *Option Pricing and Volatility: Advanced Trading Strategies and Techniques*, 2nd ed. McGraw-Hill

Rubinstein, M. (1994), "Implied binomial trees," *Journal of Finance*, 49, pp771–818

Schönbucher, P.J. (1999), "A market model of stochastic implied volatility," *Philosophical Transactions of the Royal Society*, Series A, 357, pp.2071–2092

Stein, E. and C. Stein (1991), "Stock price distributions with stochastic volatilities: An analytical approach," *Review of Financial Studies*, 4, pp727–752

List of Figures

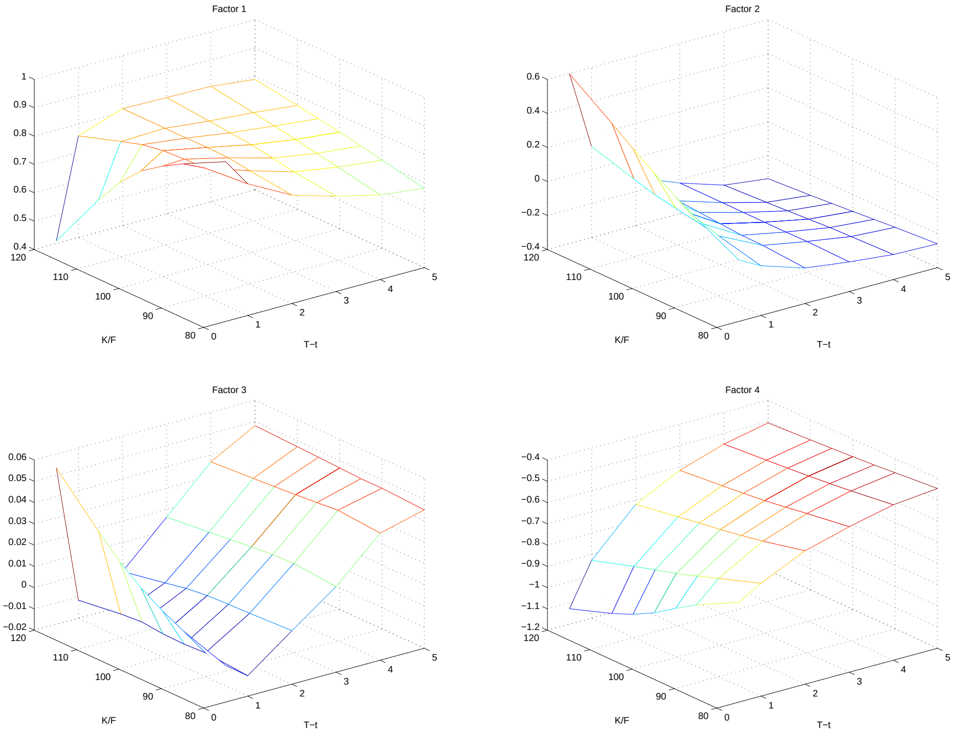

Figure 1: Estimated factors for the Totem data. The height of the surface measures the value of $\theta_{(T-t),(F/K)}$, for each maturity and moneyness combination.

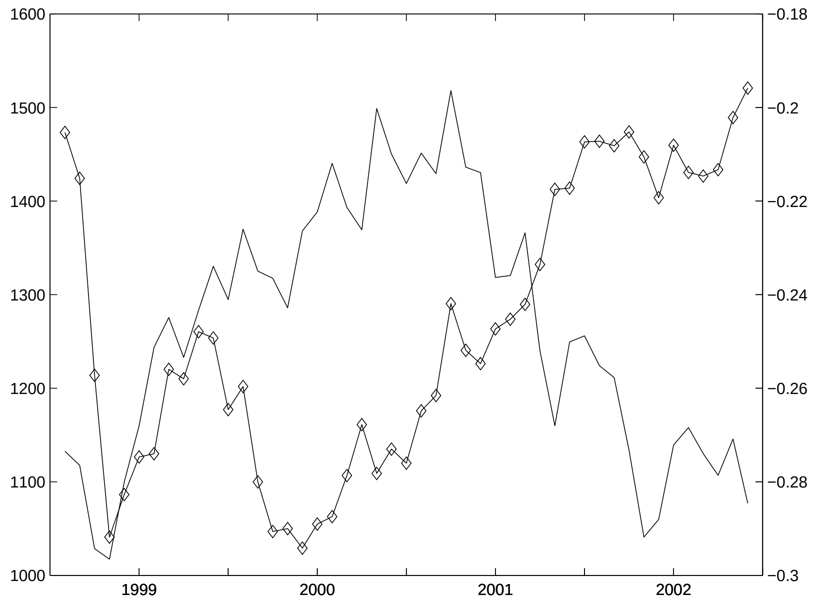

Figure 2: Plot of the level of the S&P 500 index (unmarked line, with scale on the left axis) and the estimates of c_1 (line with diamond symbols, with scale on the right axis) - our measure of skew in the volatility surface. Note the positive correlations between the two series.

List of Tables

	Time to Maturity (months)					
Strike	6	12	24	36	48	60
120	15.91%	18.49%	20.20%	21.03%	21.44%	21.64%
110	18.13%	20.33%	21.48%	21.98%	22.18%	22.30%
105	19.50%	21.26%	22.21%	22.52%	22.59%	22.70%
100	20.94%	22.38%	23.06%	23.14%	23.11%	23.07%
95	22.73%	23.71%	23.92%	23.73%	23.58%	23.53%
90	24.63%	24.99%	24.78%	24.40%	24.10%	23.96%
80	28.41%	27.71%	26.66%	25.83%	25.23%	24.83%

Table 1: Volatility matrix for June 1998. Time to maturity is measured in months while strike is in percentage terms, relative to the level of the S&P 500 index.

	Time to Maturity (months)					
Strike	6	12	24	36	48	60
0.80	0.964	0.864	0.782	0.736	0.699	0.680
0.90	0.888	0.851	0.800	0.754	0.728	0.710
0.95	0.845	0.848	0.811	0.768	0.742	0.726
1.00	0.793	0.844	0.813	0.780	0.757	0.740
1.05	0.718	0.831	0.812	0.788	0.770	0.754
1.10	0.623	0.808	0.812	0.796	0.783	0.766
1.20	0.410	0.759	0.812	0.809	0.807	0.789

Table 2: Theta values corresponding to the first factor $\rho = 1.0$. Time to maturity is measured in months while strike is in percentage terms, relative to the level of the S&P 500 index forward prices (i.e., K/F).

	Time to Maturity (months)					
Strike	6	12	24	36	48	60
0.80	-0.038	-0.107	-0.191	-0.226	-0.252	-0.260
0.90	0.091	-0.046	-0.173	-0.220	-0.261	-0.277
0.95	0.163	-0.023	-0.170	-0.221	-0.266	-0.288
1.00	0.244	-0.001	-0.160	-0.222	-0.274	-0.299
1.05	0.330	0.023	-0.150	-0.221	-0.282	-0.309
1.10	0.422	0.059	-0.140	-0.221	-0.289	-0.318
1.20	0.598	0.136	-0.119	-0.220	-0.303	-0.335

Table 3: Theta values corresponding to the second factor $\rho = 0.965$.

	Time to Maturity (months)					
Strike	6	12	24	36	48	60
0.80	-0.003	-0.010	0.005	0.020	0.040	0.045
0.90	0.004	-0.009	0.004	0.022	0.041	0.046
0.95	0.009	-0.009	0.004	0.022	0.040	0.046
1.00	0.015	-0.009	0.003	0.021	0.040	0.046
1.05	0.023	-0.008	0.002	0.019	0.039	0.047
1.10	0.033	-0.009	0.000	0.018	0.038	0.047
1.20	0.053	-0.012	-0.004	0.016	0.037	0.048

Table 4: Theta values corresponding to the third factor $\rho = 0.369$.

	Time to Maturity (months)					
Strike	6	12	24	36	48	60
0.80	-0.731	-0.670	-0.573	-0.516	-0.469	-0.448
0.90	-0.835	-0.738	-0.619	-0.545	-0.495	-0.466
0.95	-0.897	-0.773	-0.639	-0.561	-0.506	-0.475
1.00	-0.963	-0.806	-0.656	-0.576	-0.517	-0.482
1.05	-1.016	-0.835	-0.671	-0.588	-0.526	-0.489
1.10	-1.058	-0.865	-0.688	-0.598	-0.534	-0.495
1.20	-1.127	-0.928	-0.721	-0.617	-0.550	-0.507

Table 5: Theta values corresponding to the fourth factor $\rho = 0.946$.

Variable	Estimate	Standard Error	t-statistic
a_0	0.4438616	0.0238337	18.62
a_1	-0.0001944	0.000036	-5.40
a_2	0.0000000	0.0000000	1.44
a_3	-0.0262681	0.0033577	-7.82
a_4	-0.0006589	0.0003454	-1.91
a_5	0.0000290	0.0000022	13.30
R^2	0.2672		
Standard error of residuals	0.0327		

Table 6: Estimates for the version of the sticky strike model in equation (11) where volatility depends on strike price and time to maturity.

Variable	Estimate	Standard Error	t-statistic
b_0	0.0058480	0.0003801	15.39
b_1	-0.2884075	0.0019565	-147.41
b_2	0.0322727	0.0067576	4.78
b_3	-0.0075740	0.0003487	-21.72
b_4	0.0015705	0.0000701	22.42
b_5	0.0414902	0.0009180	45.20
R^2	0.9493		
Standard error of residuals	0.0057		

Table 7: Estimates for the version of the relative sticky delta model in equation (12).

Variable	Estimate	Standard Error	t-statistic
c_1	-0.2486061	0.0010002	-248.56
c_2	0.0023856	0.0031146	0.77
R^2	0.9712		
Standard error of residuals	0.0054		

Table 8: Estimates for the stationary square root of time rule in equation (13).

Test	Statistic
Equation (11) vs equation (10)	32.65
Equation (12) vs equation (11)	1.11
Equation (13) vs equation (12)	2.01

Table 9: Comparison of Models using the ratio of sums of squared errors statistic. In a two-tailed test the statistic must be greater than 1.123 (1.111) to reach the conclusion that the first equation to provide a better explanation of the data than the second equation at the 1% (5%) level