FX OPTION PRICING: RESULTS FROM BLACK SCHOLES, LOCAL VOL, QUASI Q-PHI AND STOCHASTIC Q-PHI MODELS

Krishnamurthy Vaidyanathan[1]

Abstract

The paper suggests a new class of models (Q-Phi) to capture the information that the market provides through the 25-Delta Strangles and 25-Delta Risk Reversals. The model is able to capture the stochastic movements of a full strike structure of implied volatilities.We argue that extracting information through this model and pricing path-dependent and non-benchmark strike options is a better methodology than using a contant implied volatility. The model can be used to price exotic options and hedge them robustly with benchmark European options.

1 Introduction

Pricing derivative products is about computing the right probabilities. For instance, for a foreign exchange vanilla option, we need to know the probability density function of the underlying exchange rate at maturity. This allows calculating the probability that spot at maturity is within a given interval, which for Call options is greater than strike and for Put options is less than strike. For a barrier option, we need to compute the joint probability that the exchange rate is within a given interval at maturity and that the spot did not touch the barrier on its path from option start date to maturity. Next we would multiply the appropriate probabilities with the payoff of the option and sum over over all possible outcomes of the exchange rate at maturity to get the expected payoff. The options value is this payoff discounted to the option start date.

We imply the probabilities from market prices, which for liquidly traded options is implied volatility. The implied volatility varies for different option strikes. The benchmark strikes that get traded in the foreign exchange options are 50-delta (more commonly known as at-the-money-forward or ATMF strike), 25-Delta Call and 25-Delta Put (also known as 75-delta Call). Figure 1 depicts a typical volatility smile quoted in the market. The market quotes ATMF volatility (or short form 'vol') , 25-Delta Strangle and 25-Delta Risk Reversal which are related to the 25-Delta Call and Put volatilities as follows:

25-Delta Strangle = (25-Delta Call vol + 25-Delta Put vol)/2 – ATMF vol (1)
25-Delta Risk Reversal = 25-Delta Call vol - 25-Delta Put vol (favouring Calls) (2)

Alternatively, if 25-Delta Put vol is higher than 25-Delta Call vol, 25-Delta Risk Reversal is said to be 'favouring Puts' and is quoted as (25-Delta Put vol - 25-Delta Call vol).

[1] SINE, CSRE Building, IIT Bombay, Powai, Mumbai – 400076. India.
E-Mail: vaidyanathan@alumni.iitk.ac.in

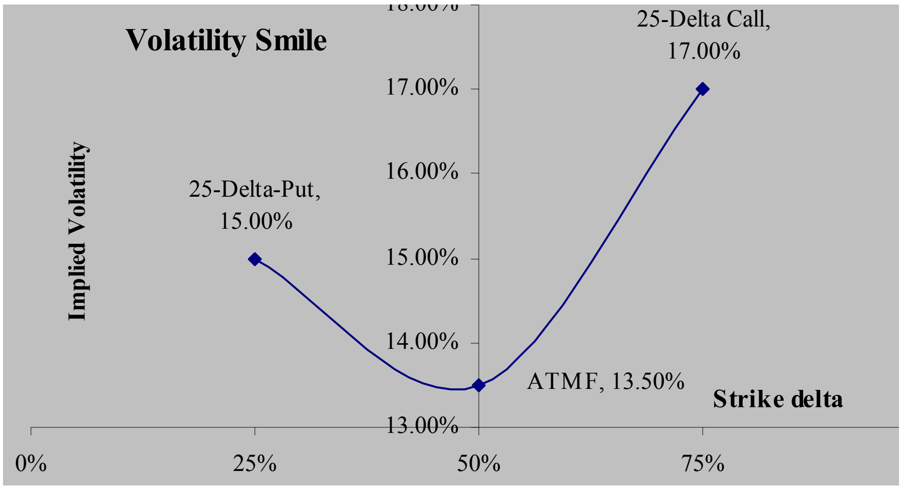

Fig. 1. The figure depicts the volatility smile that traders quote for liquid foreign exchange options. The benchmark strikes correspond to 25-Delta-Call, ATMF i.e. 50-Delta and 25-Delta-Put.

1.1 Problem of multiple volatility for same underlying

European options are often priced and hedged using the Black-Scholes (BS) model. In BS model there is a one-to one relation between the price of a European option and the volatility parameter σ_{BS} . Consequently, option prices are often quoted by stating the implied volatility σ_{BS} , the unique value of the volatility which yields the option's dollar price when used in BS. In theory, the volatility σ_{BS} in BS model is a constant. In practice, options with different strikes K require different volatilities σ_{BS} to match their market prices as shown in figure 1. Handling these market skews and smiles correctly is critical to foreign exchange desks, since they usually have large exposures across a wide range of strikes. Yet the inherent contradiction of using different volatilities for different options makes it difficult to successfully manage these risks using BS model.

1.2 The information content of benchmark options prices.

There are good reasons for incorporating the prices of benchmark option prices like 25-Delta Strangles and 25-Delta Risk Reversals into a model for pricing and risk managing foreign exchange options. Since the advent of the famous Black and Scholes (1973) option pricing model and the introduction of foreign exchange option contracts, the volume and liquidity of fx options has increased exponentially. Simultaneously more and more complex, exotic option specifications have arisen with features ranging from knock-in and knock-out barriers, digital options and range binaries to combinations of these and other features with many different payoff functions.

While on the one end of the spectrum the development has gone towards increasingly complex specifications, there has been a significant increase in liquidity in the markets of standard European Call and Put options. For almost every exchange rate, there are liquid markets for European options with a broad range of maturities and strike prices – in particular the strikes corresponding to 25-Deltas. These benchmark options make the trading of a new piece of information possible - information on volatility.

The option prices contain fundamental information on not just the volatility but also other information like the co-movement of volatility with the underlying. For instance, when the market puts a higher premium on 25-Delta Puts vis-à-vis 25-Delta Calls, it is implicitly stating that it expects the volatility to be higher when spot goes lower than the case when spot goes higher. Therefore a high risk-reversal implies a high correlation between Spot and Volatility. Similarly, a high strangle implies that the market expects volatility itself to be volatile. As an aside, if volatilities were not volatile, the options market may not exist!

Given that the benchmark option prices reveal additional information about the likely dynamics of the underlying, a pricing model should use their prices as input. This should yield an increase in accuracy over the standard Black-Scholes model. Then the standard options can also be used as additional hedge instruments.

The model presented here tries to take these points into account. It is designed to incorporate benchmark option prices and thus the information that they contain, in order to improve the pricing of more exotic instruments.

1.3 Related Literature

The deviation of observed market prices for options from their theoretical counterparts as given by the Black-Scholes formula has triggered a large literature in which both academics and practitioners alike have tried to improve on the limitations of the Black Scholes model.

One strand of the literature concentrates on the implied tree approach. The aim is to keep as closely as possible to the Black-Scholes setup while exactly reproducing the option prices given in the market. This is achieved by specifying a time and state dependent volatility function which does not contain any additional random component. Models of this type are by Rubinstein (1995), Derman and Kani (1994), Derman et.al. (1996) and Dupire (1994).

While exactly reproducing the option prices observed in the market the implied volatility models have the drawback that they do not allow for idiosyncratic stochastic dynamics in the option prices. This is in conflict with empirical observation and with the continuous updating of the new information reflected in the option prices. The poor results in a hedging test performed by Dumas et.al. (1999) are probably also due to this drawback. Dupire (1996) took a first step towards incorporating stochastic dynamics into the term structure of volatilities, but again he models realized volatilities and forward contracts on it, and not implied volatilities from options prices.

Derman and Kani (1998) extended their implied tree approach to allow for stochastic dynamics in the full term and strike structure of implied local volatilities. They derive restrictions on the drift of the local volatilities that are necessary for absence of arbitrage, and these restrictions involve integrals over all possible underlying prices and times before the maturity of the forward volatility concerned. The complexity of these restrictions makes the model hard to handle and we are going to propose a slightly different approach. Furthermore it is not obvious how in Derman and Kani's model it is ensured that the implied volatilities satisfy certain no-arbitrage restrictions as expiry is approached. The fundamental problem is, that Derman and Kani specify two things that may be contradictory: the dynamics of the spot volatility and the implied volatilities for different strike prices and maturities.

The development of local volatility models by Dupire (1994), and Derman-Kani (1996) was a major advance in handling smiles and skews. Local volatility models are self-consistent, arbitrage-free, and can be calibrated to precisely match observed market smiles and skews. Currently these models are still used for managing smile and skew risk. However, the dynamic behavior of smiles and skews predicted by local vol models is exactly opposite of the behavior observed in the marketplace: when the price of the underlying asset decreases, local vol models predict that the smile shifts to higher prices; when the price increases, these models predict that the smile shifts to lower prices. In reality, exchange rates and market smiles move in the same direction. This contradiction between the model and the marketplace tends to de-stabilize the delta and vega hedges derived from local volatility models, and often these hedges perform worse than the naive Black-Scholes' hedges.

We will follow an approach for modeling fundamental quantities like the stochastic process of the volatility of the underlying as in the traditional stochastic volatility models of Hull (1987), Heston (1993) or Stein and Stein (1991). This facilitates the fitting of the model to observed option prices and gives the model a larger degree of flexibility.

Using a model based approach means that we do not use the market-based approach applied to the term structure of implied volatilities which is similar to the market models of the term structure of interest rates by Miltersen, Sandmann and Sondermann (1995), Brace, Gatarek and Musiela (1997) and Jamshidian (1997). Nor do we model the instantaneous conditional forward volatilities as in the effective volatility model by Derman and Kani (1998), or forward variances like in Dupire (1992) or model the Black-Scholes implied volatilities.

Another direction of research has been on the nature of the underlying price process which was assumed to be a lognormal Brownian motion by Black and Scholes. Well known papers of this approach are by Hull (1987), Heston (1993), or Stein and Stein (1991) as discussed earlier.

The approach taken here draws upon both the strands of research thus far. We model the instantaneous volatility of exchange rates rather than implied volatilities as in the market models. We also incorporate stochasticity in the underlying volatility i.e. volatility of the underlying is assumed to be uncertain as is observed in the market.

Some of the constraints of our model are that though they reproduce the typical shapes of implied volatilities observed in the markets known as the smile, they need to be calibrated to the benchmark options, i.e. they are computationally intensive and there is no closed form solution for calculating option price. Another constraint is that this model has only one additional factor driving the stochastic volatility and cannot be extended to the multi-factor case.

In what follows we briefly present the models that are currently used in foreign exchange options: the famous Black-Scholes model, the Q-Φ model and the Stochastic Q-Φ model, and compare the results obtained from each of these models.

2 The Black-Scholes Model

We first introduce the Black-Scholes Model

2.1 Spot Dynamics and the Local Volatility

The Black-Scholes (1973) approach towards the probability density is straightforward. The basic assumption is that the exchange rate follows a lognormal stochastic process:

$$\frac{dS_t}{S_t} = \mu * dt + \sigma * dW_t \tag{3}$$

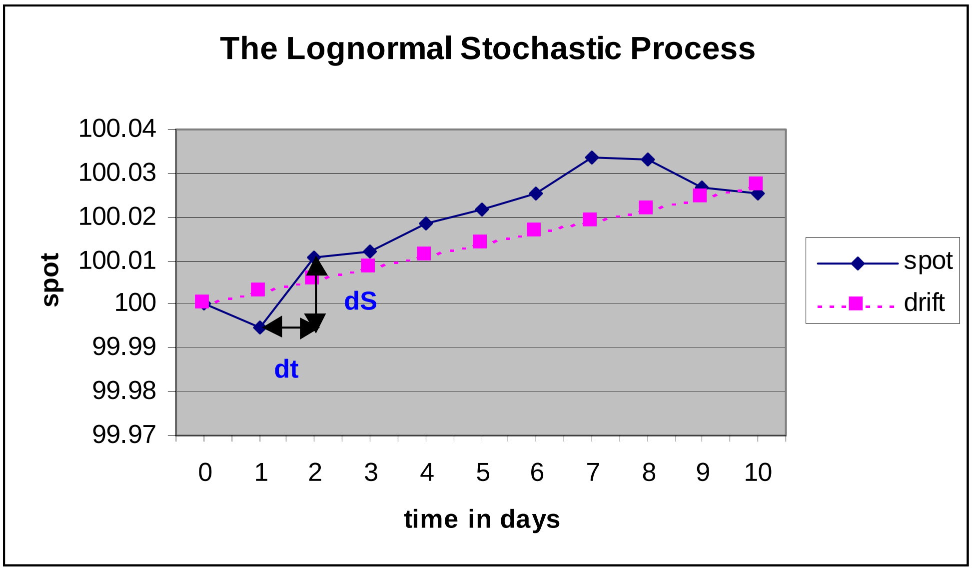

Fig. 2. The Lognormal Stochastic Process. The blue line shows a simulated path for the first 10 days of an exchange rate according to equation (3). Initial spot S_0= 100, volatility σ = 10% and drift μ = 1%. The pink line shows the drift only, i.e. the deterministic part of equation (3). The double-arrows indicate the size of the spot move dS in the second time step.

The relative move of spot dS_t/S_t over the next time interval dt is given by a deterministic part, the so-called drift term $\mu*dt$ and a random term $\sigma*dW_t$. The drift term is given by the difference of the numeraire and asset interest rates: $\mu*dt = (r_{num} - r_{ass})*dt$. The random part is driven by $dW_t = \varepsilon * \sqrt{dt}$, where ε is a random number that is drawn from a standard normal distribution $N(0,1)$. The model parameter that governs the dynamics of the Black-Scholes model is the local or instantaneous volatility σ .

We are interested in the distribution of spot at maturity of the option. So we perform a Monte Carlo simulation to create it. Draw a random number ε from our standard normal distribution $N(0,1)$, plug it into equation (3) and determine the new spot $S_{t+1} = S_t + dS_t$, which is the starting point for the next time interval dt. If we repeat this procedure until we reach the expiry date of the option we generate a single random path of spot from today's date to expiry date. A knockout barrier is taken into account by stopping any path that touches this level. The information of the outcome of S_t for many paths created in the same way provides us with a frequency distribution of outcomes of spot at expiry. After normalisation, i.e. division by the number of launched paths, we obtain the probability distribution of spot at expiry.

2.2 The Implied or Black-Scholes Volatility

In a Black-Scholes world, instead of running a tedious Monte Carlo simulation we can directly solve the stochastic partial differential equation (3)

$$S_t = S_0 * \exp\left(\left(\mu - \frac{\sigma_{BS}^2}{2}\right)t + \sigma_{BS} W_t\right) \tag{4}$$

Equation (4) allows jumping a macroscopic time step *t*, e.g. *t* = expiry date, rather than only microscopic steps of the size *dt*. σ_{BS} is the implied or Black-Scholes volatility. As ε is from *N(0,1)*, $\sigma_{BS} W_t = \sigma_{BS} \varepsilon \sqrt{t}$ is distributed $N(0, \sigma_{BS}^2 t)$, i.e. normally distributed with standard deviation $\sigma_{BS}\sqrt{t}$ and zero mean. Taking the logarithm of equation (4) shows why the Black-Scholes model is also known as the lognormal model.

$$\ln\left(\frac{S_t}{S_0}\right) = \left(\mu - \frac{\sigma_{BS}^2}{2}\right)t + \sigma W_t \tag{5}$$

Equation (5) tells us that $\ln(S_t / S_0)$ is normally distributed with mean $\mu - \sigma_{BS}^2 / 2$ and standard deviation $\sigma_{BS}\sqrt{t}$, i.e. $N(\mu - \sigma_{BS}^2/2, \sigma_{BS}^2 t)$. A simple transformation results in the PDF of S_t . The Black-Scholes or implied volatility σ_{BS} is a tenor volatility whereas σ is only valid for the next infinitesimal time step. In the Black-Scholes world these two volatilities numerically agree, i.e. $\sigma = \sigma_{BS}$, however, these volatilities are two different parameters as will become more evident when we consider more sophisticated models.

The derivation of the PDF of S_t given that the spot did not touch the barrier on its path is a bit more cumbersome but the point to be made is that we arrive at a probability density function for spot at maturity using a single model parameter σ.

2.3 Calibration

The calibration of the Black-Scholes model is simple. We determine σ_{BS} (=σ) such that the Black-Scholes price of an option is equal to the market price of this option. The standard practice in the market is to choose ATMF (at-the-money-forward) and 25-Delta Call and Put strike options for calibration.

2.4 Pricing

Pricing in the Black-Scholes world is simple too. For most option products we can provide an analytical formula for the integral of the product of probability times payoff. This makes implementation of the algorithms fast and robust.

2.5 No Smile

For a given tenor the volatility smile is defined as the implied volatility as a function of strike. Thus, per construction, the Black-Scholes model does not exhibit any smile, i.e., the function $\sigma_{BS}(K) = \sigma_{BS}$ is constant.

3 The Q-Φ Model

3.1 Spot Dynamics and the Local Volatility

The Q-Φ model allows for interest rates and hence the drift term to be time dependent. More importantly, however, the local volatility is assumed to depend on spot and time.

$$\frac{dS_t}{S_t} = \mu(t) * dt + \sigma(S,t) * dW_t \tag{6}$$

Some Wall Street houses incorporate this temporal information in order to price path-dependent options. However, the dependence of implied volatility on the strike, for a given maturity known as the Smile effect, is trickier. Many researchers have attempted to enrich the Black-Scholes model to compute a theoretical smile. Unfortunately they have to introduce a non-traded source of risk, jumps in the case of Merton (1976) and stochastic volatility in the case of Hull and White (1987), thus losing the completeness of the model. Completeness is of high value since it allows for arbitrage pricing and hedging.

If we look carefully, the process in the equation above looks very similar to the lognormal process of the Black-Scholes model that we introduced in equation (3). Yet as we simply allow the local volatility to depend on spot and time we are able to capture today's volatility smile. For option pricing we could again recourse to the Monte Carlo technique as described before for the Black-Scholes model. Draw a random number ε from $N(0,1)$, plug it into equation (6) and determine the new spot $S_{t+1} = S_t + dS_t$. However, to do so, we require the values of the local volatility for each attainable time and spot level, the so-called local volatility surface.

3.2 The Local Volatility Surface

The Q-Φ model does not provide an analytic function for σ(S,t). Rather we can derive an equation that relates the local volatility to the derivatives of the option prices with respect to strike and time.

$$\sigma^2(K,T) = \frac{r_{ass}(T) * C - \frac{\partial C}{\partial T} + \mu(T) * K * \frac{\partial C}{\partial K}}{\frac{1}{2} * K^2 * \frac{\partial^2 C}{\partial K^2}} \tag{7}$$

where C denotes the price of a Call option with tenor T and strike K. If we knew option prices for all tenors and strikes, we could numerically work out the derivatives in equation (7) and determine the local volatilities. In fact, this is the route that is taken.

Obviously, we need to know how the Q-Φ model computes vanilla option prices. They are computed as the average of two Black-Scholes prices.

$$C = \frac{1}{2}\left(BS(\sigma_1) + BS(\sigma_2)\right) \tag{8}$$

BS( ) denotes the Black-Scholes option pricing formula and the two implied volatilities are given by

$$\sigma_{1/2} = \sigma_{ATMF}(T) * \left(1 \mp Q(T) - \Phi(T) * \ln\left(\frac{F}{K}\right)\right) \tag{9}$$

Equation (8) and (9) can be interpreted in a simple way: The Q-Φ model assumes a two-point distribution for volatility. For a given strike K the volatility will be either σ_1 or σ_2 with probability 1/2. The meaning of the model parameters Q and Φ is explained later.

3.3 Calibration

We calibrate the model against the market prices for three options, ATMF, 25-delta Calls and Puts. We have to fit the three parameters of the model, σ_{ATMF}, Q, and Φ, such that the Q-Φ prices i.e., equation (8) match the market prices. Using equations (8) and (9) yields vanilla option prices for all tenors and strikes and equation (7) provides us with the volatility surface.

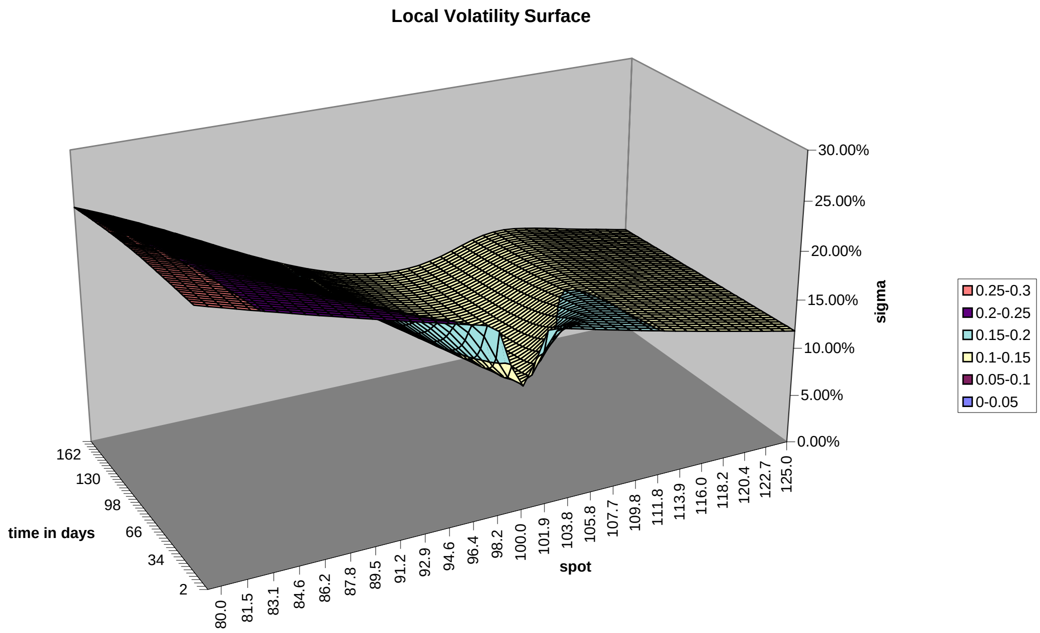

Fig. 3. Q-Φ Local Volatility Surface as a Function of Time and Spot. The graph above is generated with the Q-Φ model. It shows a 6M local volatility surface. We use an initial spot of S_0 = 100, 6M ATMF volatility σ_{ATMF} = 13.5%, strangle Str = 0.5% and risk reversal RR = 2% for Puts. The calibration returned Q = 0.36 and Φ = -1.21.

3.4 Pricing

The volatility surface is fed into a Monte Carlo simulation or a finite difference grid to price exotic options consistent with the vanilla market. Apart from vanilla options, for which we can recourse to equation (8) there is no analytical solution for Q-Φ option prices.

3.5 The Use of the Wrong Number

The implied volatility is the parameter that we have to plug into the Black-Scholes formula to obtain the market price of an option. For the Q-Φ model there is no obvious connection between the local volatility and this implied volatility. The local volatility surface contains the full information about today's smile. We can use this surface to correctly and consistently price any vanilla or exotic option. The implied volatility, in contrast, is simply the wrong number to be put into the wrong formula to get the vanilla prices right. However, using the knowingly wrong Black-Scholes formula allows market participants to express the smile in a standardised way in two single numbers: the strangle and the risk reversal.

3.6 Smile

We can easily relate Q and Φ to strangle and risk reversal. As we will show Φ introduces some asymmetry to the model while Q creates a symmetric smile.

A positive Φ tends to increase both, σ_1 and σ_2 and for options with K>F, i.e. OTM calls, and decreases both for options with K>F. Thus a positive Φ corresponds to a risk reversal that favours calls. Obviously, Φ does not have any impact on the price of ATMF options as $\ln(F/K)$ in equation (9) vanishes in that case.

Equation (8) and (9) show that *Q* does not affect the price of an ATMF option while it does so for an OTM option. Look at an option's Vega to understand this result. Vega as a function of volatility is fairly constant for ATMF options, i.e. ATMF options have zero Vomma[2]. In other words, the price of an ATMF option is an almost linear function of volatility and we have

ATMF: $$C = 1/2 * (BS(\sigma_1) + BS(\sigma_2)) \approx BS(1/2 * (\sigma_1 + \sigma_2)) = BS(\sigma_{ATMF})$$

In contrast, for OTM options Vega increases with rising volatility, i.e. they have positive Vomma. Alternatively, we say that the price of an OTM option is a convex function[3] of volatility. Thus for

OTM: $$C = 1/2 * (BS(\sigma_1) + BS(\sigma_2)) \geq BS(1/2 * (\sigma_1 + \sigma_2))$$

Apparently, *Q* makes OTM options more expensive, that is, *Q* can be related to the strangle.

[2] Vomma is the second derivative of the options price *C* with respect to volatility $\partial^2 C / \partial \sigma^2$, it describes the curvature of the options price as a function of volatility.

[3] For a convex function *c* we have $(c(x) + c(y))/2 \geq c((x + y)/2)$

4 The Stochastic Q-Φ model

We provide the formulation and dynamics of the stochastic Q-Φ model.

4.1 Spot and Volatility Dynamics

The Stochastic Q-Φ model assumes the following dynamics for the exchange rate.

$$\frac{dS_t}{S_t} = \mu(t) * dt + \sigma_t * f(S_{t,} t) * dW_t \tag{10a}$$

$$\frac{d\sigma_t}{\sigma_t} = \xi(t) * dZ_t \tag{10b}$$

$$where \quad f(S_{t,} t) = \left(\left(\frac{S}{S_0} \right)^{\alpha(t)} + 100 * \beta(t) * \left(\left(\frac{S}{S_0} \right) - 1 \right)^2 \right)$$

$$and \quad dZ_t = \kappa * \sqrt{dt} \quad with \quad \kappa \propto N(0,1)$$

The random term for the exchange rate in equation (10a) is similar to the one that we have seen in the Q-Φ model. Here, the analytical function $\sigma_t * f(S_{t,} t)$ determines the local volatility surface. However, in the stochastic model the volatility itself is a random process with no drift and a volatility (of volatility) ξ.

The model has four parameters: σ_0, α, β, and ξ. σ_0 is the initial value for the volatility. We will understand the meaning of the other parameters in what follows.

We use the Monte Carlo approach to investigate on the dynamics of equations (10a) and (10b). We first draw a random number ε, plug it into equation (10a) to determine the spot after the first time step $S_{t+1} = S_t + dS_t$. Before continuing with the next time step, however, we have to draw a random number κ, plug it into equation (10b) to obtain the correct parameter $\sigma_{t+1} = \sigma_t + d\sigma_t$ to be used for the second time step and so on. As far as the local volatility surface is concerned, the Q-Φ model is fully deterministic. In contrast, in the stochastic Q-Φ model we can only predict the local volatility for an arbitrary point in the future in a probabilistic sense. We could say that the stochastic volatility process in equation (10b) vibrates the deterministic function $\sigma_t * f(S_{t,} t)$ in a random manner. Note that this isn't exact in a mathematical sense but it is a useful way of looking at the mechanics.

We approach the full complexity of the model by first considering two special cases of the stochastic Q-Φ model, the Quasi Q-Φ model $(\xi = 0)$ and the purely stochastic model $(\beta = 0)$.

4.2 The Quasi Q-Φ Model

Suppose that the volatility of volatility vanishes $(\xi = 0)$. This means that volatility itself is not a random variable. In this case the smile is fully deterministic. The local volatility surface is given by $\sigma_t * f(S_{t,} t)$. The model becomes similar to the

Q-Φ approach (Compare equations (6) and (10a)). The strangle is reflected in β which creates a parabolic, symmetric local volatility surface around spot. The asymmetry of the smile, i.e. the risk reversal, is created by α. We have to calibrate σ_0, α, and β such that the market prices of three options, ATMF, 25Δ Calls and Puts, are matched.

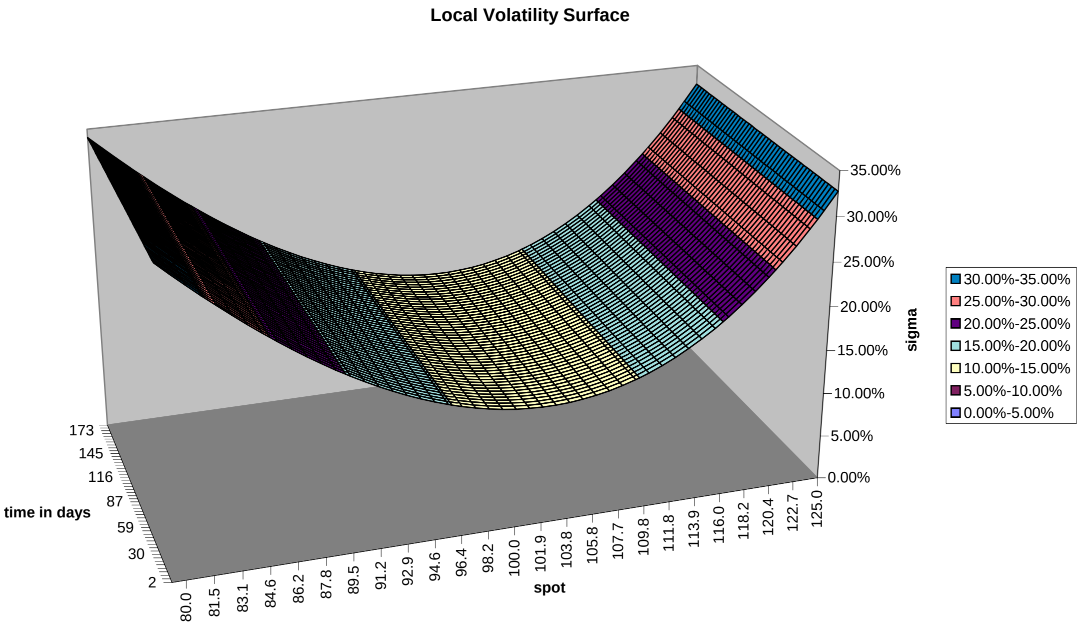

Fig. 4. Stochastic Q-Φ Local Volatility Surface. The graph above is generated with the stochastic Q-Φ model with vanishing volatility of volatility $(\xi = 0)$. For the example we used an initial spot of $S_0 = 100$, 6M ATMF volatility $\sigma_{ATMF} = 13.5\%$, strangle $Str = 0.5\%$ and $RR = 2\%$ for Puts. The calibration returned $\sigma_0 = 12.67\%$, $\beta = 0.27$ and $\alpha = -1.35$.

4.3 Q-Φ versus Quasi Q-Φ model

Compare the graphs of the local volatility in figures 3 and 4. Both surfaces were fitted to the same smile. In spite of the fact that they appear to be very different there is no contradiction here. The local volatility surface is not directly observable in the market. Rather the surface is implied and interpolated from three market observations only, the prices of a 6M ATMF option and a 6M 25Δ Call and Put. In fact the two models agree on the prices of these three options using the respective local volatility surfaces. This is a trivial result as this is simply an inversion of the calibration process. To understand in which case the models would give different results, consider the earlier Monte Carlo simulation exercise. Each path starts at $t = 0$, i.e. at the front of figure 3 or 4 and runs in a random zigzag line to the back of the graph. Along its way we pick up local volatilities and apply them to calculate the next random shock for the spot. Consider a one-touch option. A path set out to run on the surface of figure 3 presumably experiences a different probability to touch the trigger than a path running on the surface of figure 4. To summarise, two different models will naturally agree on the prices of the vanilla products that are used for the calibration. But they may give different answers, albeit the difference may be small, for products with different payoff functions, in particular path dependent options, such as knockouts and one touch options.

4.4 *Pure Stochastic Model*

Assume next that $\beta = 0$ but $\xi > 0$. In this case the risk reversal is still created by α. The strangle, however, is created by the stochastic nature of the volatility. In our simplified picture of the stochastic model the local volatility vibrates with an intensity proportional to the volatility of volatility. Thus, products with positive Vomma, such as OTM options, will become more expensive. Accordingly, the deterministic volatility surface has a nonzero slope to account for the risk reversal but does not exhibit any curvature (see figure 5). We have to calibrate σ_0, α, and ξ to the usual set of market prices: ATMF, 25Δ Calls and Puts.

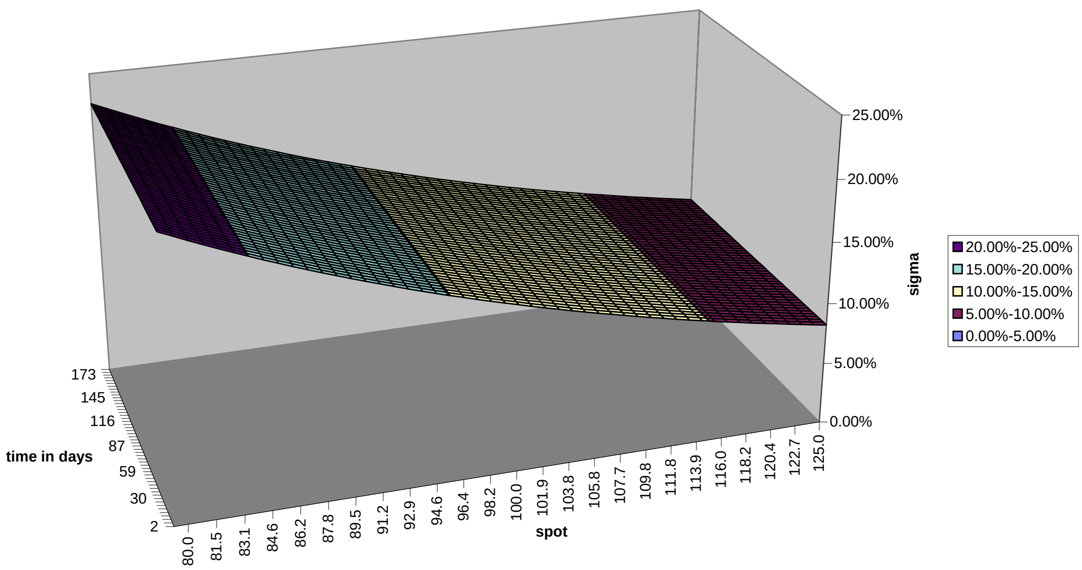

Fig. 5. Stochastic Q-Φ Local Volatility Surface. The graph above is generated with the stochastic Q-Φ model with vanishing β with similar parameters as before. The calibration returned $\sigma_0 = 12.68\%$, and $\alpha = -2.22$, and $\xi = 93\%$.

4.5 *Purely Stochastic versus Quasi Q-Φ*

What are the differences between the purely stochastic model and the Quasi Q-Φ model. To work one out, we put ourself at the starting point of a Monte Carlo simulation, i.e. at $t = 0$ and $S_0 = 100$. Suppose that after a short period of time, say, a couple of days, the random shocks bring spot down to 95. We consider the view on the volatility surface ahead with spot at 95. For the surfaces in figure 4 the landscape will be tilted, i.e. the Quasi Q-Φ model increases the risk reversal when spot moves down. In contrast, the slope of the surface in figure 5 remains unchanged, i.e. the stochastic model predicts that the risk reversal does not change at all.

Another major difference is due to the nature of the purely stochastic model. As mentioned above, due to the stochastic nature of volatility the model will mark up all positive Vomma products. In contrast, the Quasi Q-Φ model assumes a deterministic, i.e. static local volatility. It cannot capture the benefit of positive Vomma. As an example, the price of a range binary will be higher if the purely stochastic model is used as compared to the Quasi Q-Φ model.

4.6 Stochastic Q-Φ or The Mix

As described above, we either generate a smile by means of a purely deterministic local volatility surface or we use a deterministic risk reversal combined with a stochastic volatility process. Both approaches are consistent with the market, i.e. ATMF, 25Δ Calls and Puts, so we cannot determine which approach is the correct one. Additional market information is required to find the most appropriate model.

A purely deterministic volatility surface entails that the risk reversal changes quickly when spot moves, whereas in a purely stochastic model the smile shifts sideways, i.e. if the risk reversal does not change at all when spot moves. Thus, if we combine the two approaches, via the 'mix ratio' we can control our model's implied change in risk reversal with a change in spot. In fact we fit the four parameters of the full model, σ_0, α, β, and ξ such that we match ATMF, 25Δ Calls and Puts and the speed of the risk reversal with a change in spot.

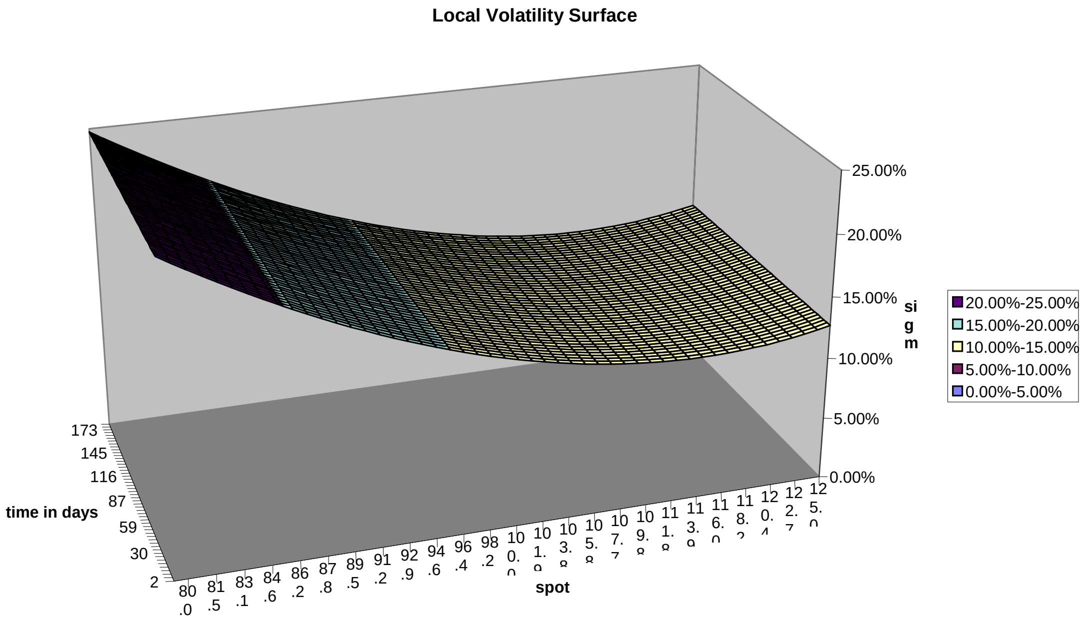

Fig. 6. Stochastic Q-Φ Local Volatility Surface. The graph above is generated with the stochastic Q-Φ model. For the example we fitted the model to an initial spot of $S_0 = 100$, 6M ATMF volatility $\sigma_{ATMF} = 13.5\%$, strangle $Str = 0.5\%$, $RR = 2\%$ for Puts, and change of risk reversal per 1% move down in spot of 0.10% The calibration returned $\sigma_0 = 12.65\%$, and $\alpha = -2.08$, $\beta = 5\%$, and $\xi = 82\%$.

β is the only model parameter that is sensitive to the speed of risk reversal. Thus, in a simplified picture, we would start the model's calibration by fitting β such that it generates the market's speed of risk reversal. At the same time, β accounts for the curvature of the local volatility surface in figure 6. However, the curvature is less pronounced as compared to the one of the surface in figure 4. This means that the strangle due to β is too small. We have to use the second handle that the model offers, the volatility of volatility ξ, to create the residual strangle. Note that ξ is smaller (ξ = 82%) as compared to the calibration result in the purely stochastic model (ξ = 93%). This is due to the fact that a part of the market strangle is already created by β.

The table below summarises the relationship of the model parameters to the market observations.

Table 1. Relationship of model parameters to at-the-money-forward options, risk-reversals, stranges and speed of risk-reversal.

	ATMF	Risk Reversal	Strangle	Speed of Risk Reversal
σ_0	+			
A		+		
B			+	+
Ξ			+	

4.7 Pricing

To price exotic options with the stochastic Q-Φ model we can again use a Monte Carlo simulation or a finite difference grid. The stochastic nature of the volatility adds another dimension to the problem. In the implementation, the algorithm must not only account for the spot moving up or down but also that the local volatility experiences random shocks. Thus, if we are to build a finite difference grid, we have to construct a cube with the axes time, spot and volatility. Note that for the Q-Φ or the Quasi Q-Φ model it is sufficient to provide a two-dimensional grid with the axes spot and time, thus making it easy to implement.

5 Pricing example

The table below summarises the prices of a $110.0/90.0\ \ 6M$ range binary. $S_0 = 100$, ATMF volatility $\sigma_{ATMF} = 13.5\%$, strangle $Str = 0.5\%$, $RR = 2\%$ for Puts, and change of risk reversal per 1% move down in spot of 0.10% using the different approaches presented above.

As the range binary exhibits a positive Vomma the price based on the purely stochastic model is highest, followed by the Stochastic Q-Φ (The Mix), which, even though smaller, still has a stochastic volatility component and finally Quasi Q-Φ, where volatility is assumed to be purely deterministic.

We observe that the speed of risk reversal which is implied by the Quasi Q-Φ model is much higher than the 0.10% that is historically seen in the market.

Table 2. Prices of range binaries for the various models considered.

	BS	Quasi Q-Φ	Purely Stochastic	Stochastic Q-Φ
Price	39.7%	42.0%	48.6%	47.5%
Speed of risk reversal	--	0.42%	0.0%	0.10%

6 Conclusions

In this paper, we propose a new class of models (Q-Φ), that captures both stochastic volatility and skewness. The models we propose are highly tractable for pricing and risk management. The model parameters are such that they can be hedged using the standard strangles and risk reversals traded in the market. The model allows for easy implementation and the pricing speed is good which is a key aspect for trading and risk management.

References

Alan Brace, Dariusz Gatarek, and Marek Musiela, 1997. The market model of interest rate dynamics. *Mathematical Finance*, 7, 127–155.

Black, F. and Scholes, M., 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81, 637-659.

Breeden, D. and R. Litzenberger, 1978. Prices of State-Contingent Claims Implicit in Option Prices. *Journal of Business*, 51, 621-651.

Derman Emanuel and Kani Iraj, 1994. Riding on the smile. *Risk*, 7, 32–29.

Derman Emanuel and Kani Iraj, 1998. Stochastic implied trees: Arbitrage pricing with stochastic term and strike structure of volatility. *International Journal of Theoretical and Applied Finance*, 1, 61–110.

Derman Emanuel, Kani Iraj, and Chriss Neil, 1996. Implied trinomial trees of the volatility smile. *Journal of Derivatives*, 4, 7–22.

Donald P. Chiras and Steven Manaster, 1978. The information content of option prices and a test of market efficiency. *Journal of Financial Economics*, 6, 213–234.

Duffie, D., 1988. *Security Markets, Stochastic Models*. Academic Press, San Diego.

Dumas Bernard, Jeff Fleming, and Robert E. Whaley (1999). Implied volatility functions: Empirical tests. *Journal of Finance*, 53, 2059–2106.

Dupire, B., 1992. Arbitrage Pricing with Stochastic Volatility. *Proceedings of AFFI Conference in Paris*, June 1992.

Dupire Bruno, 1994. Pricing with a smile. *Risk*, 7, 18–20.

Eugene F. Fama, 1970. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25, 383–420.

Farshid Jamshidian, 1997. LIBOR and swap market models and measures. *Finance and Stochastics*, Springer, 1, 293-330.

Jeff Fleming, Barbara Ostdiek, and Robert E. Whaley, 1995. Predicting stock market volatility: A new measure. *Journal of Futures Markets*, 15, 265–302.

Karatzas, I., Shreve S.E., 1988. *Brownian Motion and Stochastic Calculus*. Springer Verlag, New York.

Kristian R. Miltersen, Klaus Sandmann, and Dieter Sondermann, 1997. Closed form solutions for term structure derivatives with log-normal interest rates. *The Journal of Finance*, 52(1), 409–430.

Harrison, J.M. and D. Kreps, 1979. Martingales and Arbitrage in Multiperiod Securities Markets. *Journal of Economic Theory*, 20, 381-408.

Harrison, J.M. and S. Pliska, 1981. Martingales and Stochastic Integrals in the Theory of Continuous Trading. *Stochastic Processes and their Applications*, 11, 215-260.

Heath David, Jarrow Robert, and Morton Andrew, 1992. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60, 77–105.

Heston L. Steven, 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6, 327–343.

Hull, J. and A. White, 1987. The Pricing of Options on Assets with Stochastic Volatilities. *The Journal of Finance*, 3, 281-300.

Hull, J. and A. White, 1990. Valuing Derivative Securities Using the Explicit Finite Difference Method. *Journal of Financial and Quantitative Analysis*, 25, 87-100.

Merton, R., 1976. Option Pricing when Underlying Stock Returns are Discontinuous. *Journal of Financial Economics*, 3, 124-144.

Rubinstein Mark, 1995. As simple as one, two, three. *Risk*, 8, 44–47.

Stein E. and Stein J., 1991. Stock price distributions with stochastic volatility: An analytic approach. *Review of Financial Studies*, 4, 727–752.