

Local Volatility Modeling of JSE Exotic Can-Do Options

Antonie Kotzé^a, Rudolf Oosthuizen^b, Edson Pindza^c

^a*Senior Research Associate, Faculty of Economic and Financial Sciences
Department of Finance and Investment Management, University of Johannesburg
PO Box 524, Aucklandpark 2006, South Africa*

^b*The JSE, One Exchange Square, Gwen Lane, Sandown, 2196, South Africa
Tel: +2711-520-7000*

^c*University of Pretoria*

Number of pages: 46

Date: 1 July 2014

Abstract

Can-Do Options are derivative products listed on the JSE's derivative exchanges — mostly equity derivative products listed on Safex and currency derivative products listed on Yield-X. These products give investors the advantages of listed derivatives with the flexibility of “over the counter” (OTC) contracts. Investors can negotiate the terms for all option contracts, choosing the type of option, underlying asset and the expiry date. Many exotic options and even exotic option structures are listed. Exotic options cannot be valued using closed-form solutions or even by numerical methods assuming constant volatility. Most exotic options on Safex and Yield-X are valued by local volatility models. Pricing under local volatility has become a field of extensive research in finance and various models are proposed in order to overcome the shortcomings of the Black-Scholes model that assumes the volatility to be constant.

In this document we discuss various topics that influence the successful construction of implied and local volatility surfaces in practice. We focus on arbitrage-free conditions, choice of calibrating functionals and selection of numerical algorithms to price options. We illustrate our methodologies by studying the local volatility surfaces of South African index and foreign exchange options. Numerical experiments are conducted using Excel and MATLAB.

Keywords: Exotic options, JSE, Can-Do Options, Implied Volatility, Local Volatility, Dupire Transforms, Gyöngy Theorem, Markov Projection

(JEL) Classification: C61, G17

Email addresses: antonie@quantonline.co.za; <http://www.quantonline.co.za> (Antonie Kotzé), rudolfo@jse.co.za (Rudolf Oosthuizen), pindzaedson@yahoo.fr (Edson Pindza)

Contents

1	Introduction	3
2	Black-Scholes Partial Differential Equations	5
2.1	Assuming Constant Volatility	5
2.2	Consistency with the Volatility Skew	6
3	Local volatility	7
3.1	Rubinstein's Thoughts	8
3.2	The Similarity with Forward Rates	9
3.3	What does it Mean in Practice?	12
3.3.1	Implied and Instantaneous Volatility	12
3.3.2	Local Volatility by Example	15
3.3.3	Sticky Local Volatility	17
3.4	Local Volatility by a Deterministic Volatility Function	19
3.5	Dupire Local Volatility	19
3.6	Local Volatility in terms of Implied Volatility	23
4	The ALSI: Exact Implementation	24
4.1	The Deterministic Implied Volatility Function	24
4.2	The ALSI Implied and Local Volatilities	27
5	Numerical Implementation of Dupire	29
5.1	Volatility Interpolation and Extrapolation	29
5.2	Safex's Implementation	31
5.3	Comparing the Implied and Local Volatilities for Dtop and Usdzar	35
5.4	Dupire in terms of Call Prices	35
6	Conclusion	36
	Appendices	38
A	Derivation local volatility formula	38
A.1	Dupire local volatility	38
A.2	Local Volatility by Implied Volatility	40
B	Gyöngy's Theorem and Markov Projection	41

1. Introduction

One of the central ideas of economic thought is that, in properly functioning markets, prices of traded goods contain valuable information that can be used to make a wide variety of economic decisions. In financial derivative markets, implied volatility is one such traded quantity. The expected future value of traded volatility plays a central role in finance theory. It is thus crucial to estimate this parameter accurately enabling meaningful financial decision making.

In finance, the volatility is defined as a variational measure of price of a given financial instrument over time. There exist many types of volatilities classified using different standards. For instance, the historical or realised volatility is a type of volatility derived from time series based on the past market prices. Implied volatility, however, is a type of volatility derived from the market — obtained from traded derivatives like options — while local or instantaneous volatility is not directly measurable from the market nor from historical data. Black and Scholes assumed volatility to be constant over the life of an option which was helpful in deriving the seminal Black-Scholes formula for option prices — remember, there were other important assumptions as well (see Hull (2012); Kotzé (2003); Black (1988)).

Practitioners and quants know that the assumption of constant volatility is wrong. The evidence to this is the long-observed pattern of implied volatilities, where at-the-money and in-the-money options tend to have lower implied volatilities than out-of-the-money options for equities. This pattern is called the *volatility skew*. Implied volatility for currency options, on the other hand, has a smile pattern where in-the-money and out-the-money options have higher implied volatilities than at-the-money options.

One explanation for this phenomenon is that, in reality, the assumed random nature of asset prices is plainly wrong. Asset prices thus do not diffuse through time in a random manner implying the probability distribution of asset prices is not log-normal. Market practitioners rectify (or modify) the Black-Scholes assumption of geometric Brownian motion by trading options with different strikes at different volatilities. The fact that volatility is not constant has ramifications when path-dependent options are priced. These products do not just depend on the variability of the underlying asset's price, but also on the variability of the volatility. Now remember, the choice of which model to use depends on the different risks involved in the option (Bouzoubaa & Osserein, 2010). Path-dependent options thus has skew dependence or they are skew-sensitive. We need to use models that capture this effect and that prompt the development of various volatility models as an alternative to the Black-Scholes model. In particular, stochastic volatility models (Heston, 1993; Hull & White, 1993), local volatility models (Coleman et al., 1999; Dupire, 1994; Derman & Kani, 1994), jump-diffusion models (Kou, 2002; Merton, 1976) and Lévy models (Madan & Seneta, 1990) amongst others were developed.

In this paper, we consider the local volatility approach to value a new range of listed exotic products called Can-Do options — this product range was launched by the Johannesburg Stock Exchange (JSE) on 8 January 2007. Can-Do options are similar to CBOE's Flex option range of products. It started out where clients wanted

to have the ability to customise key contract terms like the expiry date. However, the suite of products quickly grew to include exotic options and structured products¹. Kotzé & Oosthuizen (2013) discuss and explain the local volatility pricing of exotic Can-Do options like Barrier options, as well as the methodologies used to determine their initial margins. Local volatility models have been in use since the 1980s although these were not known by the name “local volatility.” The mathematical framework for local volatility was first formulated by Dupire (1994). At the same time, Derman & Kani (1994) and Rubinstein (1994) solved this problem numerically by implementing binomial trees. These methods have subsequently been improved by many other researchers (Andersen & Andreasen, 2000; Lagnado & Osher, 1997). It has since been realised that Dupire’s framework is an extension of research done by Gyöngy (1986).

The local volatility functional approach has become popular with practitioners, because of its simplicity and the fact that it conveniently retains the market completeness² of the Black-Scholes model and it is consistent³. Due to this the Black-Scholes model is a convenient translation mechanism. It is universally used by traders to “talk” to one another; it gives unambiguous answers. It further retains the intuitive concept of the implied volatility skew where the implied volatility is a function of the strike and time only — it is non-stochastic. Similar to this parameterisation, the local volatility is also a function of strike and time. Moreover, the existence of a forward equation that describes the evolution of call option prices as functions of maturity time and strike price makes it possible to express the unknown volatility function directly in terms of known option prices. It thus captures the implied volatility skew without introducing additional sources of risk. This is different to stochastic volatility models which assume the volatility follows a random process as well.

Some local volatility parameterisations have one practical problem that needs to be addressed. In the real market, there exist a limited number of accessible data points (traded volatilities or volatility bids and offers). When implementing the local volatility models, one might have to estimate values between two consecutive given data points. This is overcome by interpolation. There are two ways this can be done: either smooth the data points by fitting a functional form (like a quadratic function) or find an interpolation technique that works (simplest being linear interpolation). However, using the incorrect interpolation technique can be dangerous in that it can yield arbitrage opportunities. Therefore it is important to consider practical arbitrage free techniques to construct volatility surfaces. This paper gives an overview of such approaches, describes characteristics of volatility surfaces and provides practical details for the construction of volatility surfaces.

The rest of this paper is organized as follows. In Section 2, we describe the formulation of the Black-Scholes partial differential equation and show how it can be

¹http://www.jse.co.za/Products/Equity-Derivatives-Market/Equity-Derivatives-Product-Detail/Can-Do_Futures_and_Options.aspx

²It allows hedging based on the underlying asset alone.

³It has no contradictions i.e., the price is the price. If two traders use the same input values, the answers will be the same, exactly.

made to incorporate the implied volatility skew. This section can be skipped if the reader do not want to re-read the mathematical rigors of the Black-Scholes theory. Section 3 reviews the concept of local volatility and explains it from a practical perspective. We introduce the Dupire formulation using Dupire’s mapping between implied and local volatility and discuss the numerical implementation thereof. Section 4 introduces the functional implementation for the Alsi implied volatility surface and we show how the Dupire framework can be used to obtain the corresponding local volatility surface algebraically. In Section 5 we show how the local volatility surface can be obtained numerically for the DTOP index and USDZAR exchange rate. We conclude in Section 6.

2. Black-Scholes Partial Differential Equations

This section discusses the mathematical background and can be skipped

We want to determine the price of a security that guarantees a payment $h(\cdot)$ contingent on the state of the process S at maturity date T . For a call option $h(S) = \max[S - K, 0]$ where K is called the strike price.

In the Black-Scholes framework, we start by assuming that the security price S evolves log-normally, according to the following stochastic differential equation (Wilmott, 1998)

$$\frac{dS}{S} = \mu dt + \sigma dW \tag{2.1}$$

where μ is the expected continuously compounded rate of return earned by an investor in a short period of time dt — the instantaneous expected return. Further, σ is the instantaneous volatility or standard deviation of return from the security price S , and W is a standard Brownian motion or Wiener process, with mean zero and variance equal to 1. A geometric Brownian motion is a natural two-parameter model of a security-price process because of the simple interpretations of μ and σ (Duffie, 1996). Note that W , and consequently its infinitesimal increment dW , represents the only source of uncertainty in the price history of the security.

2.1. Assuming Constant Volatility

Black, Scholes and Merton made some assumptions in order to facilitate a better understanding of the dynamics of the security price S . One of the main assumptions is that of risk neutrality. In its simplest form, this infers that all risk-free portfolios can be assumed to earn the same risk-free rate. They further assumed that the volatility in equation (2.1) is deterministic (constant) and the discount rate is the constant risk-free rate r . Under these assumptions, the risk-neutral dynamic of the asset is (Hull, 2012)

$$dS_t = (r - d)S_t dt + \sigma(K, T)S_t dW_t. \tag{2.2}$$

where W_t is a standard Brownian motion or Wiener process, S_t denotes a risky underlying asset price process at time t and d is the constant dividend rate. equation (2.2)

describes a simple one-factor asset price process where $\sigma(K, T)$ is called the implied volatility and it is a function of the fixed strike K and expiry time T only.

Let the scalar function $V_{bs}(S, t)$ be the value of a contingent claim like an option at any time t conditional on the price of the underlying being S at that time. Using Ito's lemma, equation (2.2) can be transformed to the Black-Scholes stochastic partial differential equation (PDE)

$$\frac{\partial V_{bs}}{\partial t} + \frac{1}{2}\sigma^2(K, T)S^2\frac{\partial^2 V_{bs}}{\partial S^2} + (r - d)S\frac{\partial V_{bs}}{\partial S} - rV_{bs} = 0. \quad (2.3)$$

equation (2.3) basically describes how the value of a derivative contract, at a continuum of potential future scenarios, diffuses *backwards* in time towards today. This equation is a backward parabolic partial differential equation also known as the backward Kolmogorov equation.

Under the assumption of a constant volatility $\sigma(K, T)$, this PDE can be solved *analytically* by applying the Feynman-Kac theorem and resulting formula (Castagna, 2010). This formula establishes a link between parabolic partial differential equations and stochastic processes. It offers a method of solving certain PDEs by simulating random paths of a stochastic process (Klebaner, 2005; Clark, 2011). The solution of equation (2.3) with the terminal condition $V_{bs}(S_T, T) = \max[S_T - K, 0]$ where S_T is the underlying's price at time T , gives the seminal Black-Scholes formula.

Please note that equation (2.3) is solved *backwards* in time — the terminal conditions are specified and the solution *today* is sought (Clark, 2011). Even though the Black-Scholes formula is the solution to a very specific and simple PDE, it is still very relevant.

One crucial point one has to understand is that the implied volatility, $\sigma(K, T)$ in equation (2.3) is not linked in any simple way to the volatility, σ , of the true stock price process described in equation (2.1). The implied volatility is not the standard deviation of S . This led to the famous statement by Rebonato (2004):

“Implied volatility is the wrong number to put into the wrong formula to get the right price of plain-vanilla options.”

Implied volatilities are simply a short-hand notation to quote a price! This implies that

$$\sigma(K, T)^2 T \neq \int_0^T \sigma(S_u, u)^2 du \quad (2.4)$$

where $\sigma(S_t, t)$ is the instantaneous volatility of the stochastic process S_t . Implied volatilities are not volatilities after all (neither instantaneous nor average)!

2.2. Consistency with the Volatility Skew

Although satisfactory for European options, the Black-Scholes model comes up short for more complex options, such as Asian options (whose payoff depends on the average price of the underlying asset over time), barrier options (whose value depends on whether a specific boundary value has been attained by the underlying asset before its maturity) or even common American options.

Practitioners thus started looking for a simple way of pricing more complex options. The prerequisite was that the methodology should be consistent with the volatility skew. If we now generalise equation (2.2) and assume that volatility is dependent on the asset's price and time (it's not constant anymore) but we still assume it to be deterministic, we get

$$dS_t = (r - d)S_t dt + \sigma(S_t, t)S_t dW_t. \quad (2.5)$$

Here, the function $\sigma(S, t)$ is called the local volatility function because it is dependent on both S and t . Note that $\sigma(t)$ is sometimes referred to as the instantaneous volatility — it is a function of time only.

Using Ito's lemma and the scalar function $V_l(S, t)$ (the value of an option), equation (2.5) can be transformed to the generalised Black-Scholes PDE

$$\frac{\partial V_l(S, t)}{\partial t} + \frac{1}{2}\sigma^2(S, t)S^2 \frac{\partial^2 V_l}{\partial S^2} + (r - d)S \frac{\partial V_l}{\partial S} - r_t V_l = 0. \quad (2.6)$$

If $V_l(S, t)$ is twice differentiable and $V_l(S_T, T)$ is the terminal condition, the Feynman-Kac theorem states that the solution to this backward parabolic partial differential equation shown in equation (2.6) is given by

$$V_l(S, t) = E^{\mathbb{Q}} \left[e^{-\int_t^T r_u du} V_l(S_T, T) | S_t = S \right], \quad (2.7)$$

where $S \in \mathbb{R}$ and S_t is described by the stochastic differential equation (2.5) and r_u is the instantaneous discount rate applicable for a very short period du (Linetsky, 1998; Duffie, 1996). Note that the expectation is taken under the risk-neutral probability measure \mathbb{Q} where the stochastic term in equation (2.5) is governed by Brownian motion or it is a Wiener process.

Black, Scholes and Merton had to assume a constant or fixed volatility to be able to obtain an analytic solution from equation (2.7). When the volatility is a function of both the spot price S and time t (even if assumed to be deterministic, not stochastic), equation (2.7) has no analytic or closed-form solution and equation (2.6) can be solved numerically only.

The big question is, what is this local volatility $\sigma(S, t)$? Can it be measured or estimated? Implied volatility $\sigma(K, T)$ is a tradable quantity and thus measurable. Local volatility is not. Equation (2.6) will only be useful if we can understand the local volatility function $\sigma(S, t)$ and if we can measure it somehow or link it to measurable quantities.

3. Local volatility

Local volatility models are widely used in the finance industry (Engelmann et al., 2009). Whereas stochastic volatility and jump-diffusion models introduce new risks into the modeling process, local volatility models stay close to the Black-Scholes theoretical framework and only introduce more flexibility to the volatility. This is one of the main reasons of fierce criticism of local volatility models (Ayache et al.,

2004). Thus, it is a mistake to interpret local volatility as a complete representation of the true stochastic process driving the underlying asset price. Local volatility is merely a simplification that is practically useful for describing a price process with non-constant volatility. A local volatility model is a special case of the more general stochastic volatility models. That is why these models are also known as “*restricted stochastic volatility models*”.

The local volatility function $\sigma(S, t)$ is assumed to be deterministic — it is a deterministic function of a stochastic quantity S_t and time. So there is still just one source of randomness, ensuring the completeness of the Black-Scholes model is preserved. Completeness is important, because it guarantees unique prices, thus arbitrage pricing and hedging (Dupire, 1993).

Dupire (1994) was the first to show algebraically that, given prices of European call or put options across all strikes and maturities, we may deduce the volatility function $\sigma(S, t)$, which produces those prices via the full Black-Scholes equation (Clark, 2011). Dupire’s insight was that if the spot price follows a risk-neutral random walk and if no-arbitrage market prices for European vanilla options are available for all strikes and expiries, then the local volatility $\sigma(S, t)$ in equation (2.5) can be extracted analytically from European option prices (Dupire, 1993). He, unknowingly, applied Gyöngy’s theorem (Gyöngy, 1986).

3.1. Rubinstein’s Thoughts

The question was how could we obtain the local volatility $\sigma(S, t)$ if we only have the traded implied volatilities $\sigma(K, T)$ at hand? What is the relationship between these two volatilities? While Dupire (1993) realised that the local volatilities should be compatible with the observed smiles at all maturities, Rubinstein (1994) reiterated the following

“One of the central ideas of economic thought is that, in properly functioning markets, prices contain valuable information that can be used to make a wide variety of economic decisions. At the simplest level, a farmer learns of increased demand (or reduced supply) for his crops by observing increases in prices, which in turn may motivate him to plant more acreage. In financial economics, for example, it has been argued that future spot interest rates, predictions of inflation, or even anticipation of turns in the business cycle, can be inferred from current bond prices. The efficacy of such inferences depends on four conditions:

- A satisfactory model that relates prices to the desired inferred information,
- A model which can be implemented by timely and low-cost methods,
- Correct measurement of the exogenous inputs required by the model, and
- The efficiency of markets.

Indeed, *given the right model, a fast and low-cost method of implementation, correctly specified inputs, and market efficiency, usually it will not be possible to obtain a superior estimate of the variable in question by any other method.*”

This is probably the reason why the simple Black-Scholes equation is still in use today — irrespective of its deficiencies.

The above mentioned thought process led Rubinstein (1994) and Derman & Kani (1994) to develop numerical schemes where they were able to relate the local volatility to the stock price, implied volatilities and time. Both of these methods use the so-called *implied trees*. The basic idea of these tree schemes is to price options in a standard Cox, Ross and Rubinstein (CRR) tree with a constant volatility, and then adjust the volatility at the nodes in the tree by using the given implied volatility skew, to obtain the correct market prices for the relevant options. The disadvantages of these methods are that they are slow and notoriously unstable while convergence seems to be a problem (Rebonato, 2004). That is why the analytic solution of Dupire (1994) is preferred over the numerical schemes. We’ll discuss this in Section 3.4.

3.2. The Similarity with Forward Rates

Let’s digress a little and introduce forward rates. Time value of money and zero coupon yield curves taught us that we can invest money today and earn interest on it. The rates we can invest at is usually the spot rates, i.e., rates that start from today and mature at a time T . A question that we need to ask is, if we know we are going to receive an amount of money at a certain point in the future, and we want to invest it at that point, can we fix the interest rate of that investment today?

The answer is yes. This is done by fixing a forward rate. A *forward rate of interest* is the rate of interest, implicit in currently quoted spot rates, that would be applicable from one time point in the future to another time point in the future (Brigo & Mercurio, 2001). In the market, a forward rate is quoted as 6X12 for instance. This points to a rate that is applicable in 6 months time and ends in 12 month’s time. A time line explanation is given in Fig. 2.

Forward rates are not usually quoted and has to be implied from spot rates through an arbitrage argument. Let’s assume we want to invest an amount A for 12 months. There are two alternative ways we can do this:

- invest at the 12 month rate given by r_2 ; or
- invest at the 6 month rate for 6 months given by r_1 . After 6 months, role the investment for another 6 months. There are risks in this due to us not knowing what the rates are going to be in 6 months time. So, let’s assume we can fix this rate today and that we reinvest at the 6X12 forward rate given by $r_{6|12}$.

Both investments mature at the same time. If there are no arbitrage opportunities, we should end up with the same amount of money. Thus (using simple rates)

$$A(1 + r_1 \times 0.5)(1 + r_{6|12} \times 0.5) = A(1 + r_2).$$




Figure 1: Determining the forward rate.

The money market spot rates for r_1 and r_2 are known and the only unknown is $r_{6|12}$. This equation can thus be rearranged to obtain $r_{6|12}$. In general we obtain

$$r_{1,2} = \frac{1}{t_{1|2}} \left[\frac{1 + r_2 t_2}{1 + r_1 t_1} - 1 \right] \quad (3.8)$$

where $r_{1,2}$ is the forward rate effective from time t_1 to t_2 in the future and $t_{1,2} = t_2 - t_1$. This rearrangement is messy for compound rates but if the rates were continuous rates we have

$$r_{1,2} = \frac{r_2 t_2 - r_1 t_1}{t_{1|2}}. \quad (3.9)$$

The explanation is given in Figure 1.

Now, in order to understand volatility's link with forward interest rates, we need to understand what significance volatility has to an option trader. Remember, historical volatility is calculated from the time series of the stock prices. This volatility is thus 'backward looking.' On the other hand, implied or market volatility is 'forward looking' i.e., it is an estimate of the future volatility or the volatility that should prevail from today until the expiry of the option.

Rational market makers base option prices on these estimates of future volatility. To them, the Black-Scholes implied volatility $\sigma(K, T)$ is, to some extent, 'the estimated average future volatility' of the underlying over the lifetime of the option. In this sense, $\sigma(K, T)$ is a *global measure of volatility*.

On the other hand, the local volatility $\sigma(S, t)$, represents 'some kind of average' over all possible *instantaneous volatilities*, at a certain point in time, in a stochastic volatility world (Gatheral, 2006). Unlike the naive implied volatility $\sigma(K, T)$ produced by applying the Black-Scholes formula to market prices, the local volatility is the volatility implied by option values produced by the one factor Black-Scholes PDE given in equation (2.6).

Derman et al. (1995) explained the analogy between trying to value a bond by picking the "correct" YTM and the dilemma in trying to value an exotic option by picking the "correct" implied volatility as follows

"In the bond market, each bond has its YTM. The YTM of a bond is actually the implied constant forward discount rate that equates the present value of a bond's coupon and principal payments to its current market price. Similarly, in the index options market, each standard option has its

own implied volatility, which is the *implied constant future local volatility* that equates the Black-Scholes value of an option to its current market price.

They also stated this differently (Kani et al., 1996)

“The forward rate from one future time to another can be found from the prices of bonds maturing at those times (see Figure 2) similarly, the local volatility at a future stock level and time is related to options expiring in that neighborhood.”

This became obvious when Dupire (1993) and Derman & Kani (1994) noted, that, knowing all European option prices merely amounts to knowing the probability densities of the underlying stock price at different times, conditional on its current value. Further insight came when they realised that under risk neutrality, there was a unique diffusion process consistent with the risk neutral probability densities derived from the prices of European options. This diffusion is unique for any particular stock price and holds for all options on that stock, irrespective of their strike level or time to expiration. Remember that under the general Black-Scholes theory, the implied volatility skew infers that one stock should have many different diffusion processes: one for every strike and time to expiry. This, of course, cannot be the case.

Kani et al. (1996) further showed that the local variance, $\sigma^2(K, T)$ is the conditional risk-neutral expectation of the instantaneous future variance of the stock returns, given that the stock's level at the future time T is K . We can also interpret this measure as a K -level, T -maturity forward-risk-adjusted measure. This is analogous to the known relationship between the forward and future spot interest rates where the forward rate is the forward-risk-adjusted expectation of the instantaneous future spot rate (Brigo & Mercurio, 2001).

They carry on

“Implied trees (or local volatility models) take a similar approach to exotic options. They avoid implied volatility, and instead use the volatility surface of liquid standard options to deduce future local volatilities. Then, they use these local volatilities to value all exotic options.”

Dupire, Derman, Kani and Kamal rediscovered a known (but lost) theorem stated and proved a decade earlier. They proofed it independently from a practitioner's point of view. This was a proposition by Krylov but proven by Gyöngy (1986). Gyöngy's theorem is an important theoretical result that links local volatility models to other diffusion models that are also capable of generating the implied volatility surface.

Alexander & Nogueira (2008) pointed out that the basic idea is that, a given stochastic differential equation (SDE) with stochastic drift and diffusion coefficients, can be mimicked by another constructed process. However, this mimicking SDE can be constructed such that it has *deterministic* coefficients such that the solutions of the two equations have the same marginal probability distributions. In essence, Gyöngy's theorem states that the local volatility SDE given in equation (2.5) is just

a special case of a more general stochastic differential equation with stochastic drift and volatility.

With Gyöngy one can map a multi-dimensional Ito process to a one-dimensional Markov process with the same marginal distributions as the original process. It is used by practitioners to construct simpler SDEs by, for instance, reducing the number of stochastic processes in models that incorporate stochastic volatility and stochastic interest rates and/or dividends. Read more on this important theorem in Appendix B.

3.3. What does it Mean in Practice?

We know that every option volatility that is traded in the market, is linked to a strike and a time as depicted by the implied volatility surface — we say traded volatility is state and space dependent. Conceptually, implied volatility represents the average amount of movement that we expect the stock to experience given that the stock price is to reach its given strike at the given time of expiration T .

Further remember that volatility measures variability, or dispersion about a central tendency — it is simply a measure of the degree of price movement in a stock, futures contract or any other market. The implied volatility is the amount that the market thinks the stock would change on average *over a period of time*. Practically, under the Black-Scholes assumption of normality, a 10% annualised volatility represents the following: in one year, returns will be within [-10%; +10%] with a probability of 68.3% (1 standard deviation from the mean); within [-20%; +20%] with a probability of 95.4% (2 standard deviations), and within [-30%; +30%] with a probability of 99.7% (3 standard deviations).

3.3.1. Implied and Instantaneous Volatility

In order to understand what is meant by “local” volatility, we need to refresh our memory on historical, implied and instantaneous volatilities. Gatheral (2006) stated that the local volatility is representing ‘some kind of average over all possible *instantaneous volatilities*’ in a stochastic volatility world. This is confusing because the implied volatility seems to be defined in exactly the same way: implied volatility is equal to the quadratic mean volatility from 0 to an expiry time T . What does this mean?

If market volatilities are not known, many traders will look at historical volatilities (also known as realised volatility). Let’s assume we have a stock price that diffuses through time. The prices are given by $S_1, S_2, S_3, \dots, S_i, \dots, S_n$ and S_i is the stock price at the end of the i -th time interval where $i = 1, 2, 3, \dots, n$. The logarithmic (per period) returns are defined as

$$x_i = \ln \left(\frac{S_{i+1}}{S_i} \right). \quad (3.10)$$

The standard deviation is then given by

$$\hat{\sigma}(T) = \sqrt{\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n}}, \quad (3.11)$$

where \bar{x} is the mean of all x_i . $\hat{\sigma}(T)$ in equation (3.11) is defined as the historical volatility over the time interval $[0, T]$. Note, this is the average volatility for our n discrete time intervals.

To enable us to compare volatilities for different interval lengths we usually express volatility in annual terms. To do this we scale this estimate with an annualisation factor (normalising constant) which is the number of intervals per annum such that

$$\sigma_{annual} = \hat{\sigma}(T) \times \sqrt{h}.$$

If daily data is used the interval is one trading day, then we use $h = 252$ business days per annum, if the interval is a week, $h = 52$ and $h = 12$ for monthly data. Note, annualisation is just a simplification of a more complex theory and cannot be used with high frequency or intra-day tick data: $\lim_{h \rightarrow \infty} \sigma_{annual} \rightarrow \infty$ explodes as h becomes larger and larger.

An example is shown in Table 1 and Figure 2 for some stock with 11 prices and 10 time intervals. If the time interval for our example was one day, then Table 1

Steps/ Time	Price (S_i)	Logarithmic Returns (x_i)	Forward Instantaneous Volatility (σ_{inst})
1	100		9.9751%
2	101	0.0099503	9.9259%
3	102	0.0098523	9.9259%
4	101	-0.0098523	20.1021%
5	97	-0.0404095	10.1274%
6	98	0.0102565	10.0759%
7	99	0.0101524	17.2780%
8	102	0.0298530	9.8773%
9	103	0.0097562	14.0030%
10	101	-0.0196085	9.9751%
11	100	-0.0099503	
	Standard Deviation	2.00%	
	Annualised Historical Volatility	31.77%	

Table 1: Prices and volatilities for some stock.

shows that the annualised volatility (or realised volatility) for this discrete time series

is 31.77%. This is just the *average* of the per period volatility of 2% multiplied by $\sqrt{252}$. This might not make sense if we have daily or even monthly data. However, these time intervals can be made as small as possible and in the limit as $n \rightarrow \infty$, we arrive at the true *instantaneous* volatility.




Figure 2: Price path of some stock showing the standard deviation, annualised volatility and instantaneous volatilities.

Figure 2 shows the prices after each time interval. Also shown is the definition of the instantaneous volatility which is just the square root of the logarithmic return per period as given by equation (3.10) or

$$\sigma_{inst}(t_i) = \sqrt{x_i}. \quad (3.12)$$

Note that instantaneous volatility is time dependent only

This definition is similar to the definition for the *instantaneous forward interest rate* eluded to in Section (3.2) (Brigo & Mercurio, 2001). In the market the smallest time interval might be one second and we can then calculate a ‘market instantaneous volatility’ from one second to the next.

It is very important to note the way we calculate the *forward* instantaneous volatility:

$$\sigma_{instant}(t = 4) = \sqrt{x_4} = \sqrt{\ln(S_5/S_4)}.$$

Furthermore, if we have a real historic time series of stock prices, we can only obtain *one* instantaneous volatility at each time step. In Table 1 we list the 10 forward instantaneous volatilities for the 11 given stock prices.

3.3.2. Local Volatility by Example

If we understand the concept of *instantaneous* volatility, we are much closer in understanding the statement by Gatheral (2006): local volatility is representing some kind of average over all possible instantaneous volatilities in a stochastic volatility world.

In Figure 2 we plotted *one* price path and from the previous section we now know we can estimate the forward instantaneous volatilities at each time step. However, instantaneous volatility is time dependent only, but we know that local volatility is state and space dependent. Instantaneous volatility alone does not, unfortunately, bring us much closer to an explanation of local volatility. Local volatility can never be obtained from market or historical stock prices. It can only be obtained through simulation or from option prices in the market. We'll discuss the algebraic mapping due to Dupire (1994) in Section 3.5. Tree numerical schemes were discussed by Derman & Kani (1994) and Rubinstein (1994).

Let's try to understand local volatility through simulation. In Figure 3 we show a typical price path of a stock as it starts at a price of R100 on 10 February 2014 and it diffuses discretely through time — this is similar to our example in Figure 2. However, these prices are simulated through a Monte Carlo process, they are not real market values. This process should be governed by the correct stochastic volatility PDE. We can obtain the forward instantaneous volatilities for each time step as explained in the previous section. From this simulation we estimate the implied/realised volatility over the 6 month period to be about 30%.




Figure 3: One simulated price path of some stock with a 30% volatility over a six month period

Let's now run another simulation and add this price path to Figure 3. This is shown in Figure 4. We can now zoom in on 26 May 2014 where the stock price level is R95. We note that for this particular example, the two simulated price paths cross

and both arrived at the same stock price at the same time. On this date, we have two price paths and we can thus calculate two forward instantaneous volatilities.



Figure 4: Two simulated price paths of some stock with a 30% volatility over a six month period

The local volatility is then defined, à la Gatheral, as the average of these two instantaneous volatilities. Why? Remember, we know that local volatility is state and space dependent, thus the specific level of the stock price has to come into play as well. Crucial is the fact that we are at an instant of time at a certain stock price.

We can generalise the simulation by generating an infinite number of price paths through simulation. Through the law of large numbers, we expect to have an infinite number of paths that will pass through the R95 level (Duffie, 1996). We can then calculate the local volatility as follows

$$\sigma(R95, 26May2014) = \sqrt{\frac{1}{m} \sum_{j=1}^m x_j},$$

where we have m ($m \rightarrow \infty$) price paths running through R95 meaning we will have m instantaneous volatilities. In other words, the “local volatility is representing some kind of average over all possible instantaneous volatilities”.

We can now generalise this further if we realise that, if we run an infinite number of simulations we will have many paths passing through every conceivable stock price S on 26 May 2014 where $S \in \mathbb{R}_0^+$. We will actually end up with many, many price paths passing through every stock price at every time interval. In practice we only have discrete time intervals and also discrete stock prices, so we can state this

mathematically as follows

$$\sigma(S_k(t_i), t_i) = \sqrt{\frac{1}{m} \sum_{j=1}^m x_j(S_k(t_i), t_i)}, \quad (3.13)$$

where t_i is the i -th time interval ($i = 1, 2, 3, \dots, n$) and S_k is the k -th price at time t_i where $k = 1, 2, 3, \dots, l$ such that we have l stock prices at each time interval.

If we do this numerically for a discrete number of stock prices and a discrete number of time intervals, we end up with a matrix where we have a local volatility for every stock price at every time interval. In Figure 4 counter i runs from left to right across time and k from bottom to top across different stock prices. In the limits as $n, m, l \rightarrow \infty$, we'll end up with the continuous version.

In our example shown in Figure 4, we can do this for every possible future time from 10 February 2014 till 10 August 2014 and every possible price at each time. If we do this we obtain the whole local volatility surface. In practise we can only do it discretely and thus we do it for every day, week or month. We can for instance divide the time to expiry into monthly buckets and choose stock levels from R10 to R200 in increments of R10. We then determine the local volatility for every price and every bucket giving us a grid of volatilities. Such a grid (matrix) is shown in Figure 5 and it represents the local volatility surface. From Figure 5 we deduce that, although local volatilities are not “real”, they look similar to the market or implied volatilities.

3.3.3. Sticky Local Volatility

We now know that local volatility is the name given for the instantaneous volatility of an underlying (i.e., the exact volatility), at a certain point in time at a specific stock level S_t . The volatility is “stuck” to the price of the underlying at that point in time and we thus call such a surface a “sticky local volatility surface.” This can be compared to a sticky strike or a sticky delta implied volatility surface.

From Sections 3.2 and 3.3.1, we know that the Black-Scholes implied volatility of an option with strike K is equal to the average (across time) local (or instantaneous) volatility of all possible paths of the underlying from spot at S_0 to strike K . This can be approximated by the average of the local volatility at spot and the local volatility at strike K . Bennett & Gil (2012) shows that this approximation leads to three results:

- The ATM Black-Scholes volatility is equal to the ATM local volatility.
- The Black-Scholes skew is half the local volatility skew (due to averaging). Skew here is similar to the slope of the curve. In the market it is usually taken as, simply, the difference between the 90% strike volatility and 100% strike volatility.
- A sticky local volatility surface implies a negative correlation between spot and implied volatility — The local volatility framework is thus very realistic (Kotzé & Joseph (2009) discusses this feature with examples).

Stock Price Levels	10	40%	38%	36%	34%	32%	30%
	20	39%	37%	35%	33%	31%	29%
	30	38%	36%	34%	32%	30%	28%
	40	37%	35%	33%	31%	29%	27%
	50	36%	34%	32%	30%	28%	26%
	60	35%	33%	31%	29%	27%	25%
	70	34%	32%	30%	28%	26%	24%
	80	33%	31%	29%	27%	25%	23%
	90	32%	30%	28%	26%	24%	22%
	100	31%	29%	27%	25%	23%	21%
	110	30%	28%	26%	24%	22%	20%
	120	29%	27%	25%	23%	21%	19%
	130	28%	26%	24%	22%	20%	18%
	140	27%	25%	23%	21%	19%	17%
	150	26%	24%	22%	20%	18%	16%
	160	25%	23%	21%	19%	17%	15%
	170	24%	22%	20%	18%	16%	14%
	180	23%	21%	19%	17%	15%	13%
	190	22%	20%	18%	16%	14%	12%
200	21%	19%	17%	15%	13%	11%	
		10-Mar-14	10-Apr-14	10-Mei-14	10-Jun-14	10-Jul-14	10-Aug-14
		Monthly Buckets					

Figure 5: A grid or matrix showing possible local volatilities per time bucket and price level.

The second point is understood through the following example:
 Let's assume the local volatility for the 90% strike is 22% and the ATM local volatility is 20%. The 90%-100% local volatility skew is therefore 2%. As the Black-Scholes 90% strike option will have an implied volatility of 21% (the average of 22% and 20%), it has a 90%-100% skew of 1% (as the ATM Black-Scholes volatility is equal to the 20% ATM local volatility). This is half the local volatility skew. The dynamics of the sticky local volatility surface is depicted in Figure 6.




Figure 6: Dynamics of a local volatility surface (Taken from Bennett & Gil (2012)).

3.4. Local Volatility by a Deterministic Volatility Function

Local volatility is not traded and thus is not a measurable quantity like implied volatility. Local volatility must be calculated somehow. This was a conundrum for a long time. Practitioners and quants knew from the start that there should be some link between the implied and local volatility. This makes sense because the implied volatility, $\sigma(K, T)$ is just a special case of the local volatility $\sigma(S, t)$ such that

$$\sigma(K, T) \propto \sigma(S, t).$$

Remember that $K, S \in \mathbb{R}_0^+$ and $t \in [0, T]$.

Before a derivative can be priced using the local volatility, we need to obtain the local volatility function. Before the mid-1990s, quants calculated the local volatility surface by pricing a vanilla option using the general Black-Scholes equation and the implied volatility from the traded skew. The local volatility could then (and still can) be found by setting

$$V_l(\sigma(S, t)) - V_{bs}(\sigma(K, T)) = 0$$

and then finding $\sigma(S, t)$ by solving equation (2.6) numerically by MC or FD methods. Optimisation is usually done by nonlinear least squares (NLS).

The local volatility surface obtained in this manner is in general not smooth. Dumas et al. (1998) realised that if volatility is a deterministic function of asset price and/or time, option valuation based on the Black-Scholes partial differential equation in (2.6) remains possible, although not by means of the Black-Scholes formula itself. They stated this is a special case and termed it the “deterministic volatility function” (DVF) hypothesis. They further stated that the reliability of this hypothesis depends critically on how well one can estimate the dynamics of the underlying asset price from a cross section of option prices.

Carr et al. (2013) recently proved that a quadratic analytic functional form for the local volatility is tractable in accordance to the local volatility SDE in equation (2.5). This was previously empirically proven by Dumas et al. (1998).

3.5. Dupire Local Volatility

Let’s recap what we know: From equation (2.5) we know that the general Black-Scholes differential equation is obtained from the SDE

$$dS_t = (r - d)S_t dt + \sigma(S_t, t)S_t dW_t. \quad (3.14)$$

In its most general form, the function $\sigma(S, t)$ is called the instantaneous volatility for an option with time to expiry T and its dynamics are stochastic. Furthermore, similar to the definition of the instantaneous forward interest rate (Brigo & Mercurio, 2001), we define the implied volatility as follows

$$\sigma(K, T) = \sqrt{\frac{1}{T} \int_0^T \sigma^2(t) dt} \quad (3.15)$$

by additivity of variance (Clark, 2011). This shows that the implied volatility is a fair average, across time, over all instantaneous variances.

Readers might be confused due to the last paragraph in Section 2.1 and equation (2.4) where we stated that the implied volatility is not a volatility after all. The instantaneous volatility/variance in equation (3.15) is not the true volatility of the stochastic process S_t . It is something different as explained in Section 3.

We further know that the Black-Scholes backward parabolic equation in variables S and t , given in (2.6), is the Feynman-Kac representation of the discounted expected value of the final option value as given in (2.7) (Cozzi, 2012).

Dupire (1993) attempted to answer the question of whether it was possible to construct a state-dependent instantaneous volatility that, when fed into the one-dimensional diffusion equation given in (3.14), will recover the entire implied volatility surface $\sigma(K, T)$? This suggests he wanted to know whether a deterministic volatility function exists that satisfies equation (3.14). The answer to this question is negative, however, it becomes true if we apply Gyöngy's theorem where the stochastic differential equation in (3.14) is transformed to another SDE with a non-random volatility function. This is ultimately what Dupire (1994) proved.

According to standard financial theory, the price at time t of a call option with strike price K , maturity time T is the discounted expectation of its payoff, under the risk-neutral measure. Dupire (1994) assumed that the probability density of the underlying asset S_t at the time t has to satisfy the forward Fokker-Planck equation (also known as the forward Kolmogorov equation) given by

$$\frac{\partial}{\partial t}\varphi = \frac{1}{2}\frac{\partial^2}{\partial S_T^2}(\sigma^2(S_T, T)S_T^2\varphi) - S_t\frac{\partial}{\partial S_T}((r-d)S_T\varphi) \quad (3.16)$$

where $\varphi \equiv \varphi(S_T, T)$ is the forward transition probability density of the random variable S_T (final S at expiry time T) in the SDE shown in equation (3.14) (Clark, 2011). However, using the the Breeden-Litzenberger formula (Breeden & Litzenberger, 1978)

$$\frac{\partial^2 C}{\partial K^2} = \varphi(K, T), \quad (3.17)$$

and doing a bit of calculus (see Clark (2011) and Gatheral (2006)), we can rewrite this and obtain the Dupire forward equation in terms of call option prices $C(K, T)$

$$\left[\frac{\partial}{\partial T} + (r-d)K\frac{\partial}{\partial K} - \frac{1}{2}\sigma^2(K, T)K^2\frac{\partial^2}{\partial K^2} + d \right] C(K, T) = 0 \quad (3.18)$$

where $\sigma(K, T)$ is continuous, twice-differentiable in strike and once in time, and local volatility is uniquely determined by the surface of call option prices. Note how, when moving from the backward Kolmogorov equation (the Black-Scholes PDE given in (2.3)) to the forward Fokker-Planck equation, the time to maturity T has replaced the calendar time t and the strike K has replaced the stock level S . The Fokker-Planck equation describes how a price propagates forward in time. This equation is usually used when one knows the distribution density at an earlier time, and one

wants to discover how this density spreads out as time progresses, given the drift and volatility of the process (Rebonato, 2004). Dupire's forward equation also provides useful insights into the inverse problem of calibrating diffusion models to observed call and put option prices.

This forward PDE is very useful because it holds in a more general context than the backward PDE: even if the (risk-neutral) dynamics of the underlying asset is not necessarily Markovian, but described by a continuous Brownian martingale

$$dS_t = S_t \sigma dW_t$$

then call options still verify a forward PDE where the diffusion coefficient is given by the local (or effective) volatility function $\sigma(S, t)$ given by

$$\sigma(S, t) = \sqrt{\mathbb{E}[\sigma^2 | S_t = S]}$$

and σ is the instantaneous or stochastic volatility of the process S_t . This method is linked to the Markovian projection problem: the construction of a Markov process which mimics the marginal distributions of a martingale. Such mimicking processes provide a method to extend the Dupire equation to non-Markovian settings (see Appendix B).

Dupire (1994) used the strike/dual Gamma ($\partial^2 C / \partial K^2$) of a call option, that gives the marginal probability distribution function (Breedon & Litzenberger, 1978), and thought of this as timeslices of the forward transition probabilities. Now, since the forward Fokker-Planck equation in equation (3.16) describes the time evolution of forward transition probabilities, we can isolate the volatility coefficients that recover the prices of tradable call options for all strikes and all times in equation (3.18).

Dupire actually proved that it is possible to find the same option price solving a dual problem, namely a *forward* parabolic equation in the variables K and T known as the dual Black-Scholes equation or Dupire's equation. This equation is actually the Fokker-Planck equation for the probability density function of the underlying asset integrated twice. This solution allows us to calculate the price of an European call option for every strike and maturity, given the present spot value S and time t .

Dupire (1994) solved for the volatility and got

$$\sigma_{loc}^2(K, T) = 2 \left\{ \frac{\frac{\partial C(K, T)}{\partial T} + C(K, T)d + (r - d)K \frac{\partial C(K, T)}{\partial K}}{K^2 \frac{\partial^2 C(K, T)}{\partial K^2}} \right\}. \quad (3.19)$$

Thus, given prices for all plain-vanilla call options $C(K, T)$ today, $\sigma_{loc}(K, T)$ is the local volatility that will prevail at time T ($t = T$) when the future stock price is equal to K ($S_T = K$). We refer the reader to Appendix A.1 for a full derivation of equation (3.19). Note that r is a fixed interest rate and d a fixed dividend yield, both in continuous format. See the number 2 on the right hand side of equation (3.19). This backs our statement in Section 3.3.3 up where we stated that the implied volatility is half the local volatility skew.

equation (3.19) ensures the existence and uniqueness of a local volatility surface which reproduces the market prices exactly

We further note that for equation (3.19) to make sense, the right hand side must be positive. How can we be sure that volatilities are not imaginary? This is guaranteed by no-arbitrage arguments. We have from the denominator

$$\sigma_{loc}^2(K, T) \propto \left(K^2 \times \frac{\partial^2 C}{\partial K^2} \right)^{-1}.$$

K^2 is positive always and, the dual gamma or risk-neutral price density must be positive in the absence of arbitrage (butterfly spread). To ensure the positivity of the numerator, we note that no-arbitrage arguments state that calendar spreads should have positive values. Kotzé & Joseph (2009) and Kotzé et al. (2013) discusses no-arbitrage arguments in detail.

An implied volatility surface is arbitrage free if the local volatility is a positive real number (not imaginary) where $\sigma_{loc}(K, T) \in \mathbb{R}_0^+$

Since, at any point in time, the value of call options with different strikes and times to maturity can be observed in the market, the local volatility is a deterministic function, even when the dynamics of the spot volatility is stochastic — Gyöngy. Also, equation (3.19) is conceptually equivalent to the Derman & Kani (1994) tree approach.

Dupire's equation (3.19) is not analytically solvable. The option price function (and the derivatives) in this equation have to be approximated numerically. The derivatives are easily obtained in using Newton's difference quotient. To obtain the dual delta, $\partial C / \partial K$, we therefore use

$$\frac{\partial C(K, T)}{\partial K} = \frac{C(K + \Delta K, \sigma(K + \Delta K, T)) - C(K, \sigma(K, T))}{\Delta K} \quad (3.20)$$

where we note that all calls should be priced with the implied volatility smile. Here, $C(K + \Delta K, \sigma(K + \Delta K, T))$ denotes a call price at a strike of $K + \Delta K$ where the implied volatility is found from the implied volatility skew for this option expiring at T with a strike of $K + \Delta K$. To find an optimal size of the bump, ΔK , it should be tested for. For an index like the Top 40, one or ten index points works very well but some people prefer it to be a percentage of the strike price. Beware not to make it too big — 1% might seem small but it is in general too big. One or 10 basis points works rather well for the time bump.

Beware, one cannot use the closed-form Black-Scholes derivatives for the dual delta. We call the numeric dual delta in equation (3.20), the impact/modified dual delta. Traders call the ordinary delta calculated this way, i.e., with the skew, the impact/modified delta due to the fact that it gives a truer reflection (or impact) of

the dynamic hedging process — it takes some of the nonlinear effects of the Black-Scholes equation, together with the shape of the volatility skew, into account. The dual gamma is numerically obtained by

$$\frac{\partial^2 C(K, T)}{\partial K^2} = \frac{C(K + \Delta K, \sigma(K + \Delta K)) - C(K - \Delta K, \sigma(K - \Delta K)) - 2C(K, \sigma(K))}{\Delta K^2}$$

where we dropped the T for clarity sake. Rebonato (2004) gives a full discussion on finding these derivatives numerically.

Solving equation (3.19) numerically has its own issues. Problems can arise when the values to be approximated are very small and small absolute errors in the approximation can lead to big relative errors, perturbing the estimated quantity heavily. Determining the density is numerically delicate. It is very small for options that are far in- or out-of-the-money (the effect is particularly large for options with short maturities). Small errors in the approximation of this derivative will get multiplied by the strike value squared resulting in big errors at these values, sometimes even giving negative values, resulting in negative variances and complex local volatilities. The local volatility will remain finite and well-behaved only if the numerator approaches zero at the right speed for these cases.

Further to the numerical issues, the continuity assumption of option prices is, of course, not very realistic. In practice option prices are known for certain discrete points and at limited number of maturities (quarterly for instance like most Safex options). The result of this is that in practice the inversion problem is ill-posed i.e. the solution is not unique and is unstable — the positivity of the second derivative in the strike direction is not guaranteed.

3.6. Local Volatility in terms of Implied Volatility

The instability of equation (3.19) forces us to consider alternatives. We know options are traded in the market on implied volatility and not price. Can we thus not transform this equation such that we supply the implied volatilities instead of option prices?

This can be done if a change of variables is made in (3.19) by rewriting the Black-Scholes equation for a call C in terms of the log-strike y where $y = \log(\text{strike}/\text{Forward})$. This leads to (Wilmott (1998) and Clark (2011))

$$\sigma_{loc}^2(K, \tau) = \frac{\sigma_{imp}^2 + 2\tau\sigma_{imp}\frac{\partial\sigma_{imp}}{\partial\tau} + 2(r-d)K\tau\sigma_{imp}\frac{\partial\sigma_{imp}}{\partial K}}{\left(1 + Kd_1\sqrt{\tau}\frac{\partial\sigma_{imp}}{\partial K}\right)^2 + K^2\tau\sigma_{imp}\left(\frac{\partial^2\sigma_{imp}}{\partial K^2} - d_1\sqrt{\tau}\left(\frac{\partial\sigma_{imp}}{\partial K}\right)^2\right)}, \quad (3.21)$$

where

$$d_1 = \frac{\ln(S_0/K) + ((r-d) + \sigma_{imp}^2/2)\tau}{\sigma_{imp}\sqrt{\tau}},$$

and $\tau = T - t$ such that t and S_0 are respectively the market date, on which the volatility smile is observed, and the asset price on that date. Note that equation

(3.21) gives the variance, i.e., σ^2 . See Appendix A.2 for the derivation.

When comparing equation (3.19) with (3.21) it is clear that the problem of numerical large errors no longer exists. The transformation of Dupire’s formula into one which depends on the implied volatility ensures that the dual Gamma is no longer alone in the denominator as it was in (3.19). The second derivative of the implied volatility is now one term of a summation, so small errors in it will not necessarily lead to large errors in the local volatility function. However, small differences in the input volatility surface can produce a big difference in the estimated local volatility.

The main problem is that the implied or traded volatilities are only known at discrete strikes K and expiries T . This is why the parameterisation chosen for the implied volatility surface is very important. If implied volatilities are used directly from the market, the derivatives in equation (3.21) needs to be obtained numerically using finite difference or other well-known techniques. This can still lead to an unstable local volatility surface. Furthermore we will have to interpolate and extrapolate the given data points unto a surface. Since obtaining the local volatility from the data involves taking derivatives, the extrapolated implied volatility surface cannot be too uneven. If it is, this unevenness will be exacerbated in the local volatility surface showing that it is not arbitrage free in these areas.

In the foreign exchange market, options are traded on the Delta — effectively a measure of the moneyness — as opposed to the absolute level of the strike. See Clark (2011) for the FX version of equation (3.21)

4. The ALSI: Exact Implementation

4.1. The Deterministic Implied Volatility Function

The issues relating to determining the derivatives in equation (3.21) are circumvented in two ways:

- Choose a particular functional form for the implied volatility surface and fit this function to the market volatility data.
- Choose a particular functional form for the local volatility surface and find it using non-linear optimisation techniques.

This is the route Safex took for the liquid ALSI options. Safex uses the following three parameter deterministic quadratic function as a good model of fit for the ALSI implied volatility data on every expiry date (for a full discussion see Kotzé & Joseph (2009) and Kotzé et al. (2013))

$$\sigma_{model}(\beta_0, \beta_1, \beta_2, T_k) = \beta_0^k + \beta_1^k M + \beta_2^k M^2. \quad (4.22)$$

T_k are the times to expiry for all listed expiry dates such that $k = 1, 2, \dots, n$ and n is the total number of listed expiries. This equation, of course, describes a parabola (for each expiry) and this conic section has the following properties (loosing the superscript k)

- $M = S/K$ is the strike price in moneyness format (Spot/Strike),
- β_0 is the constant volatility (shift or trend or level) parameter, $\beta_0 > 0$. Note that

$$\sigma \xrightarrow{M \rightarrow 0} \beta_0.$$

This means that β_0 is similar to the Y -intercept for a parabola. We have, in the Cartesian plane, the Y -axis being σ_{model} and the X -axis being the moneyness.

- β_1 is the correlation (slope) term. This parameter accounts for the negative correlation between the underlying index and volatility. The no-spread-arbitrage condition requires that $-1 < \beta_1 < 0$.
- β_2 is the volatility of volatility ('vol of vol' or curvature/convexity) parameter. The no-calendar-spread arbitrage convexity condition requires that $\beta_2 > 0$.

Note that equation (4.22) is linear in the wings if $M \rightarrow \pm\infty$ and it holds for discrete expiry time T . These properties and constraints are the factors describing the shape of the skews — they are not the no-arbitrage constraints. They are, however, related to the no-arbitrage conditions imposed on a volatility skew and surface — these are discussed in Kotzé et al. (2013).

The n parabolas described by equation (4.22) can now be fitted to the relevant option data (obtained per expiry date) independently to give us n discrete skews. In order to incorporate the time dependence and generate a continuous smooth implied volatility surface we also need a specification or functional form for the *at-the-money* (ATM) volatility term structure. It is, however, important to remember that the ATM volatility is intricately part of the skew. This infers that the two optimisations (one for the skews and the other for the ATM volatilities) can not be done strickly separate from one another. Taking the ATM term structure together with each skew will give us the 3D implied volatility surface.

It is well-known that volatility is mean reverting; when volatility is high (low) the volatility term structure is downward (upward) sloping Gatheral (2006). Safex uses the following exponential functional form for the ATM volatility term structure

$$\sigma_{atm}(M, t) = \frac{\theta}{t^\lambda}. \quad (4.23)$$

Here we have (See Kotzé & Joseph (2009) for full details)

- t is the time to expiry,
- λ controls the overall slope of the ATM term structure; $\lambda > 0$ implies a downward sloping ATM volatility term structure, whilst a $\lambda < 0$ implies an upward sloping ATM volatility term structure, and
- θ controls the short term ATM curvature.

We now assume that each of the parameters β in equation (4.22) must adhere to equation (4.23) such that

$$\beta_i(t) = \frac{\theta_i}{t^{\lambda_i}}; \quad i = 0, 1, 2. \quad (4.24)$$

To ensure the surface is continuous we dropped the superscript k and changed the discrete expiry times T_k to a continuous time t where $T_k \in t$; $t \in \mathbb{R}_0^+$. Combining Equations (4.22) and (4.24) leads to

$$\sigma_{imp}(M, t) = \frac{\theta_0}{t^{\lambda_0}} + \frac{\theta_1}{t^{\lambda_1}}M + \frac{\theta_2}{t^{\lambda_2}}M^2. \quad (4.25)$$

equation (4.25) fully describes the 3D implied volatility surface. We have to find 6 parameters by fitting this function to the market volatility skews using optimisation techniques. From equation (4.23) we know that λ controls the slope of the curves (we also call this the scaling parameters) while θ controls the curvature. From equation (4.22) we thus have

- β_0 and λ_0 are called the level or shift or trend parameters and they control the level of the volatility.
- β_1 and λ_1 are called the rho or tilt parameters and they control the overall slope of the curves.
- β_2 and λ_2 are the volatility of volatility parameters and they control the curvature.

Kotzé & Joseph (2009) mentions they use a floating surface such that $M_{ATM} = 1$ i.e. the ATM strike is one. A floating surface is defined such that

$$\sigma_{imp}(M, t) = \sigma_{ATM}(t) + \sigma_{imp}^{float}(M, t) \quad (4.26)$$

where $\sigma_{ATM}(t)$ is the market at-the-money volatility.

Now, from equation (4.25) we have

$$\sigma_{imp}^{ATM}(t) = \frac{\theta_0}{t^{\lambda_0}} + \frac{\theta_1}{t^{\lambda_1}} + \frac{\theta_2}{t^{\lambda_2}}. \quad (4.27)$$

Thus, combining equations (4.25), (4.26) and (4.27) leads us to the final 3 dimensional market volatility surface given by (transforming back from moneyness to the real world i.e. substituting $M = K/S$)

$$\sigma_{imp}(S, K, t) = \sigma_{ATM}(t) + \frac{\theta_1}{t^{\lambda_1}} \left(\frac{K}{S} - 1 \right) + \frac{\theta_2}{t^{\lambda_2}} \left(\left(\frac{K}{S} \right)^2 - 1 \right). \quad (4.28)$$

Safex obtains the ATM volatilities from the market meaning $\sigma_{ATM}(t)$ is a constant for every t .

Safex publishes two other parameters: θ_A and λ_A . These parameters are the at-the-money parameters from equation (4.23) and gives the theoretical ATM term structure of volatility. They are obtained by fitting equation (4.23) to the market (or mark-to-market (MtM)) term structure of ATM volatilities $\sigma_{MtM}^{ATM}(t)$ such that

$$\sigma_{model}^{ATM}(t) = \frac{\theta_A}{t^{\lambda_A}} \simeq \sigma_{MtM}^{ATM}(t). \quad (4.29)$$

Using them in calculating the ATM volatilities will give slightly different values if compared to the at-the-money volatilities calculated using equation (4.27). The reason is that θ_A and λ_A are obtained by fitting equation (4.23) to the raw ATM volatility data while the other θ 's and λ 's are obtained by optimising the whole surface and by taking the no-arbitrage conditions into account (see Kotzé & Joseph (2009)).

We need to mention a practical implementation point here. The term structure of ATM volatilities as obtained from the model in equation (4.29) will not coincide with all the traded or mark-to-market ATM volatilities due to the numerical fitting procedure. We must, however, ensure that if we price an option expiring on a particular date, that $\sigma_{ATM}(t)$ in equation (4.28) equates the market ATM volatility for that date. As an example, if we price a 9 month option ($T = 0.75$) and we have the 9 month mark-to-market ATM volatility $\sigma_{MtM}^{ATM}(0.75)$, we need to ensure that $\sigma_{ATM}(0.75)$ in equation (4.28) is equal to this volatility. This is achieved by floating $\sigma_{model}^{ATM}(t)$ up or down by a constant amount such that $\sigma_{ATM}(0.75) = \sigma_{MtM}^{ATM}(0.75) = \sigma_{model}^{ATM}(0.75)$. This will in general have the effect that $\sigma_{ATM}(t)$ is not equal to the mark-to-market volatilities for $t \neq 0.75$.

The whole volatility surface is now described by a functional form given in equation (4.28). The derivatives in Dupire's local volatility function in equation (3.21) can now be obtained analytically such that we have from equation (4.28)

$$\begin{aligned} \frac{\partial \sigma_{imp}(S, K, t)}{\partial K} &= \frac{1}{S} \frac{\theta_1}{t^{\lambda_1}} + 2 \frac{K}{S^2} \frac{\theta_2}{t^{\lambda_2}} \\ \frac{\partial^2 \sigma_{imp}(S, K, t)}{\partial K^2} &= 2 \frac{1}{S^2} \frac{\theta_2}{t^{\lambda_2}} \\ \frac{\partial \sigma_{imp}(S, K, t)}{\partial t} &= -\lambda_1 \theta_1 t^{-(\lambda_1+1)} \left(\frac{K}{S} - 1 \right) - \lambda_2 \theta_2 t^{-(\lambda_2+1)} \left(\left(\frac{K}{S} \right)^2 - 1 \right). \end{aligned} \quad (4.30)$$

Using Equations (4.30) will lead to a smooth Dupire local volatility function for $\sigma_{loc}(K, T)$ in equation (3.21). Further, if the implied volatility surface in equation (4.28) is arbitrage-free, the corresponding Dupire local volatility surface should be arbitrage-free as well. See the discussion on this in Section 3.5.

equation (4.28) and (4.30) can now be substituted into Dupire's equation in (3.21) such that we calculate the local volatility surface algebraically across time and strike.

4.2. The ALSI Implied and Local Volatilities

It is now pretty straight forward to obtain the local volatility surface for Alsi options. We need the four parameters, θ_1 , θ_2 , λ_1 and λ_2 and the published ATM

volatilities. These parameters are published every two weeks when Safex updates their volatility surfaces.

In Table 2 we show the parameter values, θ_i and λ_i , ($i = 1, 2, 3, ATM$) as published by Safex on 28 May 2014. Also shown are the values for θ_i/t^{λ_i} , $i = 0, 1, 2, 3$. Table 3 lists σ_{imp}^{ATM} , the model ATM volatilities and official Safex ATM volatilities for all expiry dates. In Figure 7 we show the Alsi implied and corresponding local volatility surfaces.

		Curvature			
		Rho(θ_1)	VolVol (θ_2)	Level (θ_0)	Atm (θ_{ATM})
	In Months	-0.8488985	0.1945430	0.9139862	0.1350075
	In Years	-0.4337514	0.1069264	0.4753064	0.1595808
		Decay			
		Rho (λ_1)	VolVol (λ_2)	Level (λ_0)	Atm (λ_{ATM})
		0.2702186	0.2408592	0.2631310	-0.0672942
Date	T	θ_1/t^{λ_1}	θ_2/t^{λ_2}	θ_0/t^{λ_1}	
19-06-2014	0.06027397	-92.655786%	21.033029%	99.531201%	
18-09-2014	0.30958904	-59.544292%	14.181881%	64.708854%	
18-12-2014	0.55890411	-50.759237%	12.301016%	55.393271%	
19-03-2015	0.80821918	-45.943944%	11.255306%	50.269616%	
18-06-2015	1.05753425	-42.724414%	10.549535%	46.836131%	
17-09-2015	1.30684932	-40.349172%	10.025150%	44.298712%	
15-12-2016	2.55342466	-33.668883%	8.531503%	37.140432%	
21-12-2017	3.56986301	-30.754183%	7.869980%	34.005870%	

Table 2: Optimised parameters for the Alsi deterministic implied volatility function on 28 May 2014 (see Equations (4.26), (4.27) and (4.28)).

Expiry Date	Expiry Time	Model ATM	Safex ATM
	T	$\sigma_{model}^{ATM}(t)$	$\sigma_{MtM}^{ATM}(t)$
19-06-2014	0.06027397	13.209622%	14.25
18-09-2014	0.30958904	14.747329%	14.00
18-12-2014	0.55890411	15.345386%	14.50
19-03-2015	0.80821918	15.731053%	15.00
18-06-2015	1.05753425	16.018262%	15.75
17-09-2015	1.30684932	16.248072%	16.75
15-12-2016	2.55342466	16.997206%	18.50
21-12-2017	3.56986301	17.384842%	21.00

Table 3: Model and official Safex ATM volatilities on 28 May 2014.

Continuing with our example in the previous section: if we want to price an

18 Dec 2014 Alsi option, we'll have to float the model ATMs down by $0.15345386 - 0.145 = 0.00845386$.

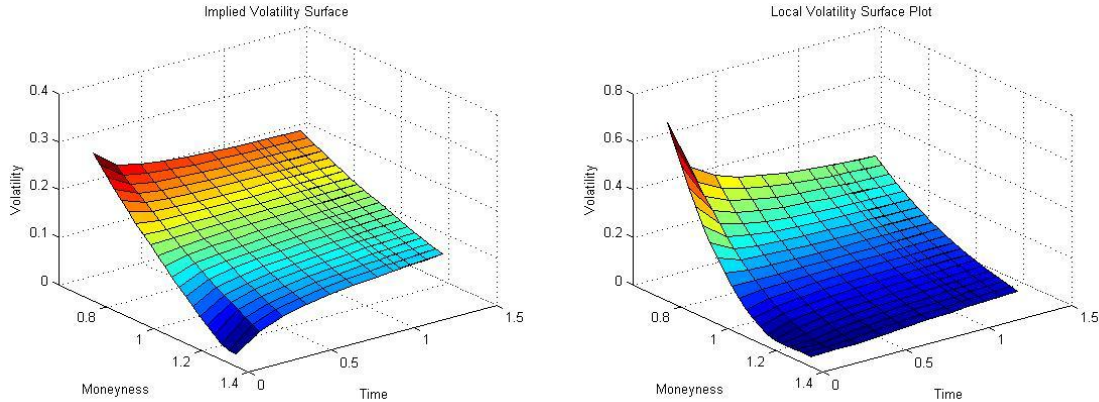


Figure 7: ALSI implied and local volatility surfaces on 28 May 2014

From Figure 7 we notice that the implied volatility surface does not have a lot of curvature — it is skewed but flat. However, we also see from the local volatility surface that it has more curvature. This shows that the local volatility skew is twice that of the implied volatility as stated in Section 3.3.3.

5. Numerical Implementation of Dupire

In Section 4.1 we described the methodology implemented for obtaining the Alsi volatility surface and how the Dupire local volatility surface can be calculated algebraically. Unfortunately, such deterministic functions for the implied volatility surfaces for some of the other instruments like the DTOP (FTSE/JSE Shareholders Weighted Top 40 Index) and USDZAR (US Dollar and South African Rand exchange rate), do not exist. The implied volatility surfaces are, however, available, albeit in a discrete form. All derivatives in equation (3.21) have to be computed numerically. The procedures and methodologies implemented at the JSE are discussed in this section.

5.1. Volatility Interpolation and Extrapolation

In practice, we are often confronted with situations where only limited amount of data is accessible and it is necessary to estimate values between two consecutive given data points. We can construct new points between known data points by interpolation or smoothing techniques.

All volatility surfaces of all derivatives listed on the JSE are available online. The following web sites should be visited:
 Equity indexes: <https://www.jse.co.za/downloadable-files?RequestNode=/Safex/mtmdata/VolSkewIndices>

Single stock futures: <https://www.jse.co.za/downloadable-files?RequestNode=/Safex/mtmdata/VolSkewSSF>
 Bonds: https://www.jse.co.za/downloadable-files?RequestNode=/YieldX/IRD_Vol_Skews

Only a handful of discrete points are given. The volatility surface (first 4 expiries) for the DTOP as published on 28 May 2014 is shown in Figure 8.

EXPIRY	19-Jun-2014	Vol	18-Sep-2014	Vol	18-Dec-2014	Vol	19-Mar-2015	Vol
LOWEST STRIKE	6850	29.47	6850	24.42	6950	23.15	7000	22.41
STRIKE	7800	23.67	7850	20.60	7950	19.97	8000	19.59
STRIKE	8800	17.99	8850	17.08	8900	17.19	9000	16.99
STRIKE	9300	15.31	9350	15.42	9400	15.81	9550	15.66
STRIKE	9750	13.00	9800	14.00	9900	14.50	10050	14.50
STRIKE	10250	10.53	10300	12.49	10400	13.25	10550	13.40
STRIKE	10750	8.17	10800	11.05	10900	12.06	11050	12.36
STRIKE	11700	4.00	11800	8.40	11900	9.88	12050	10.44
HIGHEST STRIKE	12700	0.03	12750	6.15	12900	7.94	13050	8.74
FUTURE PRICE	9750		9800		9900		10050	
BASE VOLATILITY	13.00		14.00		14.50		14.50	
MAX VOLATILITY	65.00		65.00		65.00		65.00	
MIN VOLATILITY	10.00		10.00		10.00		10.00	

Figure 8: DTOP volatility surface as published by the JSE on 28 May 2014.

Inter- and extrapolation is thus necessary. At the JSE, we take the volatilities, square them and do linear inter- and extrapolation on the variance. However, it is a two-dimensional problem. We have to do this across strike and across time. If you just need to interpolate across strike prices use the following

$$\sigma^2(K) = \sigma_1^2 + (\sigma_2^2 - \sigma_1^2) \frac{K - K_1}{K_2 - K_1}.$$

Here we want to find the variance for a certain strike given by K where $K_1 \leq K \leq K_2$ and the volatility at K_1 is σ_1 and that at K_2 is σ_2 . Further, $\sigma_1^2 \leq \sigma^2 \leq \sigma_2^2$.

When we interpolate across time only, we use what is known as *flat forward interpolation*. Volatility is time dependent. Suppose that N time points $t_0 = 0, t_1, t_2, \dots, t_N$ are given, and the implied volatilities for each of these time points are $\sigma_1, \sigma_2, \dots, \sigma_N$. We now define a general interpolation scheme to find the volatility at any time point t where $t \in \mathbb{R}_0^+$ such that (Clark, 2011)

$$\sigma(t) = \begin{cases} \sigma_1, & t < t_1 \\ \sqrt{\frac{1}{t} [\sigma_i^2 t_i + \sigma_{i,i+1}^2 (t - t_i)]}, & t_i \leq t < t_{i+1} \text{ for } i < N \\ \sigma_N, & t \geq t_N \end{cases} \quad (5.31)$$

where

$$\sigma_{i,i+1}^2 = \frac{\sigma_{i+1}^2 t_{i+1} - \sigma_i^2 t_i}{t_{i+1} - t_i}.$$

5.2. Safex's Implementation

We use the following 3 steps when we want to implement Dupire's equation in (3.21) for all instruments except the ALSI

1. Read in the implied volatility surface as published by the JSE and convert it to a floating and moneyness format;
2. Regularise the surface, meaning we interpolate and plot it with more than the given 9 strikes per expiry;
3. Use this regularised implied volatility surface when we transform it to the local volatility surface.

The first step encompass the format of how the skew is read into our model. All volatility surfaces are given in the format as shown in Figure 8. This is converted to a floating surface where the strikes are given in terms of moneyness. Moneyness is (K/S) and is just the percentage of how far the strike is in- or out-the-money. This is shown in Table 4.

DTOM4	19-Jun-14	DTOU4	18-Sep-14	DTOZ4	18-Dec-14
70.2600%	16.4700%	69.9000%	10.4200%	70.2000%	8.6500%
80.0000%	10.6700%	80.1000%	6.6000%	80.3000%	5.4700%
90.2600%	4.9900%	90.3100%	3.0800%	89.9000%	2.6900%
95.3800%	2.3100%	95.4100%	1.4200%	94.9500%	1.3100%
100.0000%	0.0000%	100.0000%	0.0000%	100.0000%	0.0000%
105.1300%	-2.4700%	105.1000%	-1.5100%	105.0500%	-1.2500%
110.2600%	-4.8300%	110.2000%	-2.9500%	110.1000%	-2.4400%
120.0000%	-9.0000%	120.4100%	-5.6000%	120.2000%	-4.6200%
130.2600%	-12.9700%	130.1000%	-7.8500%	130.3000%	-6.5600%

Table 4: Floating DTOP skews for 3 expiry times on 28 May 2014.

In Table 4 the first row contains the contract codes and skew dates. The first column gives the strikes in moneyness format. The second columns give the floating or relative volatilities. These are the volatilities relative to the at-the-money (ATM) volatilities. As an example, if the future level was 9,898 (we call this the at-the-money level) for the 18-Dec-14 expiry, we would describe the ATM volatility as the fair volatility to trade an option with a strike of 9,898. We would then expect the 10,898 ($9,898 \times 110.1\%$) strike to trade at a volatility of -2.44% relative to the ATM volatility or 2.44% below the ATM volatility. If the ATM volatility was 14.5%, the fair value volatility for this option with a strike of 10,898 would be 12.06%.

Crucial: we need the ATM volatilities. These are published by Safex daily. We show the data for the DTOP as published on 28 May 2014 in Table 5. The second last column is empty because none of the ATM volatilities changed from the previous day.

The reason for using floating skews is the fact that the SAFEX systems, \mathcal{N} utron and \mathcal{A} uclears, use them in this way. Equity skews are updated every second week

					VOL.	PREV.
DTOP	SPOT	BID	OFFER	M-T-M	CHANGES	VOLS.
19-Jun-2014	9727	9757	9757	9757		13.00
18-Sep-2014	9727	9807	9807	9807		14.00
18-Dec-2014	9727	9898	9898	9898		14.50
19-Mar-2015	9727	10015	10015	10015		14.50
18-Jun-2015	9727	10059	10059	10059		14.50
17-Dec-2015	9727	10249	10249	10249		14.50

Table 5: DTOP mark-to-market (MtM) data on 28 May 2014.

only. Between updates, the skews “float” up and down stuck to the ATM volatility at the 100% strike. The ATM volatilities are published, and these might change, on a daily basis. Changes are shown in the column “Changes” (second last) in Table 5. If this is empty, use the last column as the current ATM volatilities.

Important note:

- When we price an exotic option, we use the theoretical forward levels and not the published futures level. The reason being that barriers, for instance, are always on the cash level and not the futures level. We thus use

$$F = Se^{-(r-d)T}$$

when calculating the forward. We thus need the following inputs

- Spot or cash level S : this is the current spot price of the underlying shown in column two in Table 5.
- The current valuation date and expiry date. We then have

$$T = \frac{\text{Expiry Date} - \text{Valuation Date}}{365}$$

which is the annual time to expiry⁴.

- ATM volatility for each expiry date. This is given in the last column of Table 5.
- EndDate: The Date that the ATM volatility is applicable for. More specifically the expiry date of the option.
- The current continuous compounding interest rate r . Obtained from the official JSE zero coupon swap yield curve.
- The current continuous compounding dividend rate d . Obtained from Bloomberg.

⁴South Africa’s day count convention is actual/365.

Let's do an exercise on how we convert our floating skews back into absolute values — these are after all the values we are going to use. On 22 May 2014 we have the following for the Dtop (see Table 5):

- Spot/cash level = 9,727
- Forward level = 9,898
- Valuation date = 28-May-14
- ATM volatility = 14.50%
- Expiry Date = 18-December-14
- $r = 6.11\%$
- $d = 2.98\%$

Using these values lead us to the volatility skews shown in Table 6.

DTOM4	19-Jun-14	DTOU4	18-Sep-14	DTOZ4	18-Dec-14
6847	30.9700%	6865	24.9200%	6949	23.1500%
7796	25.1700%	7867	21.1000%	7948	19.9700%
8796	19.4900%	8870	17.5800%	8899	17.1900%
9295	16.8100%	9371	15.9200%	9399	15.8100%
9745	14.5000%	9822	14.5000%	9898	14.5000%
10245	12.0300%	10322	12.9900%	10398	13.2500%
10745	9.6700%	10823	11.5500%	10898	12.0600%
11694	5.5000%	11826	8.9000%	11898	9.8800%
12694	1.5300%	12778	6.6500%	12898	7.9400%

Table 6: DTOP implied volatility skews in absolute terms, using inputs as shown.

Remember, we need the variance or volatility squared. This is shown in Table 7 for our example. This is the end of step 1.

In step 2, we regularise our skews. This is necessary because, as shown in Tables 4, 6 and 7, we have 9 strikes only (per expiry) and these are not equidistant. To create a finer grid of strikes, with a regularised or equidistant spacing, we take the distance between the minimum and maximum strikes (as given) and divide that up into 30 intervals per expiry. This will give a grid with 31 points per expiry on the Y -axis. In our example shown in Table 7, we see the maximum strike is 12,898 and the minimum is 6,847. The grid intervals are then given by

$$((12898 - 6847)/30) = 201.7.$$

The grid will then start at 6,847 in the top left hand corner (on 19 June 2014) and end at 12,898 at the bottom right hand corner (on 18 Dec 2014) with increments of 201.7. On the X -axis of our grid, we have the times to expiry — this remains as is.

DTOM4	19-Jun-14	DTOU4	18-Sep-14	DTOZ4	18-Dec-14
6847	9.5914%	6865	6.2101%	6949	5.3592%
7796	6.3353%	7867	4.4521%	7948	3.9880%
8796	3.7986%	8870	3.0906%	8899	2.9550%
9295	2.8258%	9371	2.5345%	9399	2.4996%
9745	2.1025%	9822	2.1025%	9898	2.1025%
10245	1.4472%	10322	1.6874%	10398	1.7556%
10745	0.9351%	10823	1.3340%	10898	1.4544%
11694	0.3025%	11826	0.7921%	11898	0.9761%
12694	0.0234%	12778	0.4422%	12898	0.6304%

Table 7: Dtop variance skews. Volatilities from Table 6 converted

Next, we need the corresponding volatilities at each grid point. This is obtained through inter- and extrapolation. In practice we use Matlab’s standard “interp1” function to interpolate and extrapolate the variance linearly for each respective time. The standard formula for this method is given in equation (5.1).

As a final precaution we test all data points to ensure we never have volatilities above 100% or below 1% — we clamp our volatilities to lie between these points. This grid forms the base for all further calculations. Finally, to enable us to do temporal interpolation we convert the relative variance (being relative to time) into a total variance, simply by multiplying the variance by the relevant time in years from start date. Why do we do this? Remember from the Black-Scholes equation that the volatility is scaled by time through the factors $\sigma\sqrt{(T)}$ and σ^2T . This is the end of step 2.

The third and final step entails the implementation of Dupire’s formula in equation (3.21). At this point we note that the given strike and expiry time might still not fall on any given grid point. This is especially the case when we calculate the partial derivatives numerically. For such cases we make use of “bilinear interpolation⁵” on the grid to arrive at a total variance for the given strike and time. This method first interpolates linearly on the Y -axis (strike) using equation (5.1) and then uses the same formula to interpolate on the X -axis (time). All values are then converted from a total variance (scaled by time) to an unscaled variance by dividing by time. These numbers are tested for our allowable variance range where $0.01 \leq \sigma^2 \leq 1$, and are then converted back to volatility.

When the derivatives in equation (3.21) need to be calculated numerically, we use Newton’s difference quotient or finite difference formula. Differentiation with respect to time is implemented where we bump the time up by 1 basis point. We thus use

$$\frac{\partial\sigma}{\partial t} \approx \frac{f(t \times (1 - h)) - f(t)}{t \times h}.$$

⁵http://en.wikipedia.org/wiki/Bilinear_interpolation

When we differentiate with strike (dual Delta) we use a two-sided estimation where

$$\frac{\partial \sigma}{\partial K} \approx \frac{f(K \times (1 + h)) - f(K \times (1 - h))}{K \times (2h)},$$

and the dual Gamma is given by

$$\frac{\partial^2 \sigma}{\partial K^2} \approx \frac{f(K \times (1 + h)) - 2f(K) + f(K \times (1 - h))}{(K \times h)^2}.$$

The optimal h is found by using different h values and looking for stability. h is usually 1, 10, 25, 50 or 100 basis points.

5.3. Comparing the Implied and Local Volatilities for Dtop and Usdzar

Let's be practical and look at the volatility surfaces for the Dtop and USDZAR on 28 May 2014. The ATM volatilities and future levels are all shown in Table 5. The volatility surfaces for the Dtop is shown in Figure 9.

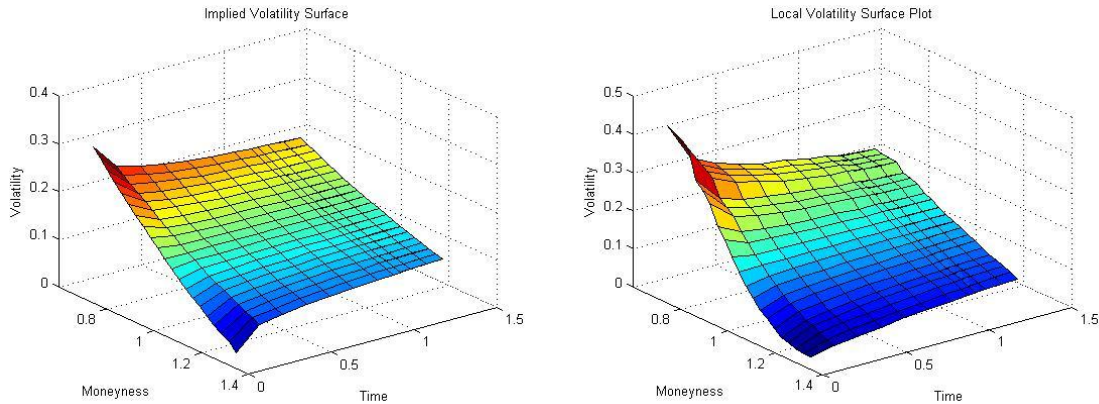


Figure 9: DTOP implied and local volatility surfaces on 28 May 2014

From Figure 9 we notice that the implied volatility surface is smooth while the local volatility surface is a bit uneven. Compare this to the smooth local volatility surface for the ALSI shown in Figure 7. The DTOP's local volatility surface is not smooth purely due to numerical differentiation. However, in Figure 10, the USDZAR implied volatility surface shows the currency market's all familiar smile. Here we also show the local volatility surface. It is still a smile but with steeper sides.

5.4. Dupire in terms of Call Prices

We can now show the instability in Dupire's local volatilities calculated in using call prices as described in equation (3.19). We implemented this equation for the Alsi local volatility surface on 28 May 2014. We did this two ways. We first used the algebraic fitted implied volatility surface using the parabola implementation in equation (4.28) where the parameters are given in Table 2. We also used the discrete Alsi volatility surface as published by Safex on this date. The two obtained local




Figure 10: USDZAR implied and local volatility surfaces on 28 May 2014

volatility surfaces are shown in Fig. (11). The instabilities are clearly seen when we are far in- and out-the-money. We have to force the volatilities to be zero when these become extremely large — this happens when the density (dual gamma) is extremely small. These plots also show that using the algebraic implementation of




Figure 11: ALSI local volatility surfaces using call prices on 28 May 2014

the implied volatility surface leads to a bit more stability, but only just! However, if these plots are compared to the Alsì's local volatility surface in Figure 7, it is clear that using Dupire's local volatility function in terms of implied volatility, equation (3.21), leads to a much more stable, practical and usable surface. Figure 9 even shows that equation (3.21) with the discrete implied volatility surface is still much more stable than using call prices.

6. Conclusion

During 2007 the JSE introduced a new class of listed derivatives. They call it Can-Do options and most of these listed options are exotic in nature. Exotic options can not be priced using the published closed-form formulae (if available). Many exotic options can, however, be priced in a local volatility framework. In this document we

discussed the local volatility framework and how the JSE implemented it ensuring correct pricing of many of its listed exotic options.

We started by introducing the general Black-Scholes-Merton stochastic differential equation. We mentioned that it can be solved analytically by using the Feynman-Kac theorem if and only if one assumes a constant volatility, interest rate and dividend yield. We further stated that this equation is also known as the Kolmogorov backward SDE because it is solved backwards in time.

We explained that implied volatility is not actually a “real” volatility because it is dependent on both time and strike. Everything was then set to introduce the concept of local volatility and explained it at hand of forward interest rates. We then gave an explanation from a practitioner’s point of view to get an intuitive feel for this mysterious concept showing how local volatilities can be obtained by simulation. Following on from this we introduced Dupire’s framework of constructing a state-dependent instantaneous volatility function that recovers the whole implied volatility surface. We discussed the forward Kolmogorov SDE and showed how to obtain the Dupire equation in terms of call prices and implied volatilities.

Safex uses a deterministic volatility functional for the implied volatility surface for Alsi options. We discussed this functional and showed how it can be implemented in the Dupire framework. We discussed the difference between the Alsi implied and local volatility surfaces. We went further by looking at the local and implied volatility surfaces for the Dtop index and USDZAR foreign exchange rate. There are no functional forms for their implied volatility surfaces so we showed how to implement Dupire numerically where we discussed how to efficiently calculate the derivatives in the Dupire equations numerically. It is shown that the Dtop implied and local surfaces look very similar to the Alsi’s (as expected). However, the USDZAR implied volatility surface has the familiar smile shape and the local volatility surface looks quite different if compared to the equity surfaces.

In the last section we showed how unstable the Dupire equation is if implemented using call option prices. This is purely due to the numerical errors when finding the derivatives numerically.

The implementation steps of Safex’s approach to implementing the local volatility surface was been set out in detail. We showed that the proposed approach proves to be simple to implement and perform well on real market data.

Acknowledgements

Thanks are due to the JSE for data and to Ania Ostaszewicz and Thorsten von Boetticher for many helpful discussions and valuable comments that helped in improving this document tremendously.

Appendices

A. Derivation local volatility formula

A.1. Dupire local volatility

In this section we derive the local volatility formula using the Dupire's model formulation. For a full derivation and explanation also see Wilmott (1998), Rebonato (2004) and Clark (2011).

Dupire sought a one-dimensional state-dependent diffusion for the underlying asset S_t of the form (see equation (2.5))

$$dS_t = \mu S_t dt + \sigma(S_t, t) S_t dW_t. \quad (\text{A.32})$$

Here, W_t is a standard Brownian motion or Wiener process, S_t denotes a risky underlying asset price process at time t , μ is the drift and $\sigma(S_t, t)$ is a local volatility (see equations (2.2) and (2.5)).

The main assumption behind the derivation of Dupire's local volatility formula is that the probability density $\varphi(S_t, t)$ of the underlying asset at the time t has to satisfy the Fokker-Planck equation such that (Dupire, 1994)

$$\frac{\partial \varphi}{\partial t} = \frac{1}{2} \frac{\partial^2}{\partial S_t^2} (\sigma^2 S_t^2 \varphi) - \frac{\partial}{\partial S_t} (\mu S_t \varphi). \quad (\text{A.33})$$

Let the discounted money market account be $F(t, T) = \exp\left(-\int_t^T r_s ds\right)$ and $C = C(S_t, K)$ the price of European options with strike K at the time t given by

$$\begin{aligned} C(S_t, K) &= F(t, T) \mathbb{E} [(S_T, K)^+], \\ &= F(t, T) \int_K^\infty \varphi(S_T, T) (S_T - K) dS_T. \end{aligned} \quad (\text{A.34})$$

The Dupire's local volatility formula is obtained by replacing equation (A.34) and its derivatives into the Fokker-Planck equation (A.33). The first derivative of (A.34) with respect to K is given by

$$\frac{\partial C(S_t, K)}{\partial K} = -F(t, T) \int_K^\infty \varphi(S_T, T) dS_T, \quad (\text{A.35})$$

and the second order derivative is

$$\begin{aligned} \frac{\partial^2 C(S_t, K)}{\partial K^2} &= -F(t, T) \frac{\partial}{\partial K} \left[\int_K^\infty \varphi(S_T, T) dS_T \right] \\ &= -F(t, T) \left[\varphi(S_T, T) \right]_{\lim_{S \rightarrow K}}^{\lim_{S \rightarrow \infty}} \\ &= F(t, T) \varphi(K, T). \end{aligned} \quad (\text{A.36})$$

Here the $\lim_{S_T \rightarrow \infty} = 0$, since we are dealing with a log normal distribution of S under geometric Brownian motion.

The first derivative of (A.34) with respect to T is given by

$$\begin{aligned} \frac{\partial C(S_T, T)}{\partial T} &= \frac{\partial}{\partial T} \left[F(t, T) \int_K^\infty \varphi(S_T, T) (S_T - K) dS_T \right] \\ &= \frac{\partial}{\partial T} [F(t, T)] \int_K^\infty \varphi(S_T, T) (S_T - K) dS_T \\ &\quad + F(t, T) \int_K^\infty (S_T - K) \frac{\partial}{\partial T} [\varphi(S_T, T)] dS_T. \end{aligned} \quad (\text{A.37})$$

Note that $\frac{\partial F}{\partial T} = -r_T F(t, T)$. Hence we can write (A.38) as

$$\frac{\partial C(S_T, K)}{\partial T} = -r_T C + F(t, T) \int_K^\infty \frac{\partial}{\partial T} \{ \varphi(S_T, T) \} (S_T - K) dS_T. \quad (\text{A.38})$$

If we substitute the Fokker-Planck equation (A.33) into (A.38) we obtain the following

$$\frac{\partial C(S_T, K)}{\partial T} = -r_T C + F(t, T) \int_K^\infty \left\{ \frac{1}{2} \frac{\partial^2}{\partial S_T^2} (\sigma^2 S_T^2 \varphi) - \frac{\partial}{\partial S_T} (\mu S_T \varphi) \right\} (S_T - K) dS_T. \quad (\text{A.39})$$

From the call price equation (A.34) and its first derivative with respect to the strike price K we have

$$\int_K^\infty S_T \varphi(S_T, T) dS_T = \frac{C}{F(t, T)} - \frac{K}{P(t, T)} \frac{\partial C}{\partial K}. \quad (\text{A.40})$$

From equation (A.37) we obtain

$$\varphi(S_T, T) = \frac{1}{F(t, T)} \frac{\partial^2 C}{\partial K^2}. \quad (\text{A.41})$$

Now, let $\mu = r_T - d_T$ where r_T is the interest rate and d_T is the dividend yield — both assumed to be deterministic. If we further substitute (A.40) and (A.41) into (A.39) and integrate by part twice, we obtain the Dupire equation given by

$$\frac{\partial C}{\partial T} = \frac{1}{2} \sigma^2 K^2 \frac{\partial^2 C}{\partial K^2} + (r_T - d_T) K \left(C - \frac{\partial C}{\partial K} \right) - r_T C. \quad (\text{A.42})$$

We can rearrange variables to obtain

$$\sigma^2(K, T) = \frac{\frac{\partial C}{\partial T} + d_T C + (r_T - d_T) K \frac{\partial C}{\partial K}}{\frac{1}{2} \sigma^2 K^2 \frac{\partial^2 C}{\partial K^2}}. \quad (\text{A.43})$$

Dupire (1994) assumed that the function $\sigma(K, T)$ is continuous, twice-differentiable in strike and once in time, and local volatility is uniquely determined by the surface of call option prices.

A.2. Local Volatility by Implied Volatility

The local volatility can be described as a function of the implied volatility if a change of variables is made in (3.19) by using C as a function of some other variable. In general, this is not possible, because a closed form formula for the call C to be transformed remains intractable. However, we can make use of the Black-Scholes formula and the concept of implied volatility. We define the following variables (Gatheral, 2006)

$$y = \ln \left(\frac{K}{F_T} \right), \quad w = T \sigma_{imp}^2(K, T), \quad \text{and} \quad F_T = S_0 \exp \left(\int_0^T \mu(t) dt \right), \quad (\text{A.44})$$

where y , w , σ_{imp} and F_T represent the log-moneyness, the total Black-Scholes implied variance, the implied volatility and the forward price, respectively. Using these change of variables the Black-Scholes call can be written as

$$C_{BS}(S_0, K, \sigma_{imp}, T) = C_{BS}(S_0, F_T e^y, w, T) = F_T [N(d_1) - e^y(d_2)], \quad (\text{A.45})$$

where

$$d_1 = \frac{\ln \left(\frac{S_0}{K} \right) + \int_0^T \mu(t) dt + \frac{w}{2}}{\sqrt{w}} = -yw^{-1/2} + \frac{1}{2}w^{-1/2},$$

$$d_2 = d_1 - \sqrt{w} = -yw^{-1/2} - \frac{1}{2}w^{-1/2}.$$

To obtain the local volatility function (3.19) in term of the market implied volatility we need to express the derivatives $\frac{\partial C}{\partial T}$, $\frac{\partial C}{\partial K}$ and $\frac{\partial^2 C}{\partial T^2}$ in terms of the new variables defined in (A.44) using the chain rule.

Following Gatheral (2006) the local volatility function $\sigma(K, T)$ is given in terms of the Black-Scholes total implied volatility w in the form

$$\sigma^2(K, T) = \frac{\frac{\partial w}{\partial T}}{1 - \frac{y}{w} \frac{\partial w}{\partial y} + \frac{1}{2} \frac{\partial^2 w}{\partial y^2} + \frac{1}{4} \left(-\frac{1}{4} - \frac{1}{w} + \frac{y^2}{w} \right) \left(\frac{\partial w}{\partial y} \right)^2} \quad (\text{A.46})$$

Alternatively, the local volatility function can be expressed in term of the implied volatility as follows

$$\sigma^2(K, T) = \frac{\sigma_{imp}^2 + 2T \sigma_{imp} \frac{\partial \sigma_{imp}}{\partial T} + 2(r_T - d_T)KT \sigma_{imp} \frac{\partial \sigma_{imp}}{\partial K}}{\left[1 + K d_1 \sqrt{T} \frac{\partial \sigma_{imp}}{\partial K} \right]^2 + K^2 T \sigma_{imp} \left[\frac{\partial^2 \sigma_{imp}}{\partial K^2} - d_1 \sqrt{T} \left(\frac{\partial \sigma_{imp}}{\partial K} \right)^2 \right]}, \quad (\text{A.47})$$

where

$$d_1 = \frac{\ln(S_0/K) + ((r_T - d_T) + \sigma_{imp}^2/2)T}{\sigma_{imp}\sqrt{T}}, \quad d_2 = d_1 - \sigma_{imp}\sqrt{T}. \quad (\text{A.48})$$

B. Gyöngy's Theorem and Markov Projection

Krylov and Gyöngy (1986) were interested in the construction of stochastic differential equations whose solutions *mimick* certain features of the solutions of Itô processes. Gyöngy showed that one can build a time-inhomogeneous Markovian process solution of an SDE that has the same one-dimensional marginals as the complicated non-Markovian Itô process itself (Atlan, 2006). Alexander & Nogueira (2008) stated this differently in that the mimicking SDE is another SDE with *deterministic coefficients* such that the solutions of the two equations have the same marginal probability distributions.

Gyöngy's Theorem is an important theoretical result because it links local volatility models to other diffusion models that are also capable of generating the implied volatility surface (Haugh, 2010). Instead of considering local volatility as a different and alternative model to stochastic volatility, it is possible to show that the first kind of model is actually a particular case of the second one — local volatility is a special case of the more general stochastic volatility.

Dupire (1994) studied the the local volatility surface that is implied, not by market prices of options, but by prices generated from a stochastic volatility model. Using infinitesimal calendar and butterfly spreads, he presents a financial argument that the square of the local volatility function is the expected value of the instantaneous squared stochastic volatility conditioned on the level of the underlying asset price. Gatheral (2006) stated this more clearly: if we suppose that the underlying asset follows a diffusion process with a stochastic instantaneous variance, then we can think of local volatility as the conditional expectation of this instantaneous volatility.

Dupire (1993) basically found that the local volatility model mimics the European option prices of some more complicated market process, and this is equivalent to matching the one-dimensional marginal distributions of that process under the equivalent martingale probability measure (also called the risk-neutral measure) used for pricing. This is exactly what Gyöngy (1986) stated and Dupire essentially recovered Gyöngy's result, albeit in a non-rigorous fashion.

We can now state Gyöngy's theorem (Alexander & Nogueira, 2008): Suppose that X_t is a real-valued one-dimensional Itô process starting at $X_0 = 0$ with dynamics

$$dX_t = \alpha(t, \omega)dt + \beta(t, \omega)dW_t \quad (\text{B.49})$$

where W_t is a k -dimensional Wiener process on the probability space $(\Omega, \mathfrak{F}, P)$, the possibly random coefficients $\alpha(t, \omega)$ and $\beta(t, \omega)$ satisfy the regularity condition of an Itô process and $\omega \in \Omega$ denotes the dependence on some arbitrary variables. In particular, if β^T is the transpose of β and we suppose that $\beta \cdot \beta^T$ is positive, then there exist another, but one-dimensional, stochastic process \tilde{X}_t that is a solution of

the stochastic differential equation

$$d\tilde{X}_t = a(t, \tilde{X}_t)dt + b(t, \tilde{X}_t)d\tilde{W}_t, \quad \tilde{X}_0 = 0 \quad (\text{B.50})$$

with *non-random* coefficients a and b defined by

$$\begin{aligned} a(t, x) &= \mathbb{E}[\alpha(t, \omega) | X_t = x] \\ b(t, x) &= \left(\mathbb{E}[\beta \cdot \beta^T(t, \omega) | X_t = x] \right)^{\frac{1}{2}}. \end{aligned} \quad (\text{B.51})$$

These deterministic coefficients then admits the same marginal probability distribution as that of X_t for every $t > 0$ and $x \in \mathbb{R}_+^d$. That is, for every Itô process of the type in equation (B.49) there is a deterministic process shown in equation (B.50) that ‘mimics’ the marginal distribution of X_t for every t .

Brunick & Shreve (2013) extended Gyöngy in two ways:

- They removed the conditions of nondegeneracy and boundedness on the covariance of the Itô process to be mimicked, requiring only integrability of this process and thereby extending the result to cover popular stochastic volatility models such as the one due to Heston (Haugh, 2010).
- They showed that the mimicking process can preserve the joint distribution of certain functionals of the Itô process (e.g., running maximum and running average) at each fixed time.

Gyöngy’s theorem leads us to the important concept of Markov Projection. When the stochastic differential equation for \tilde{X}_t in equation (B.50) has a unique solution, this solution is Markov. We thus call \tilde{X}_t the Markov projection of X_t .

Markov Projection is a very general and powerful approach in deriving accurate approximations of European-style option prices, that is, to estimate $\mathbb{E}[(S(T) - K)^+]$ when the stock price $S(t)$ follows a complicated SDE. Pitarberg (2007) summarises it in four steps

1. For the underlying of interest, its SDE, driven by a single Brownian motion, is written down by calculating its quadratic variance (and combining all dt terms that exist).
2. In this SDE, the diffusion and drift coefficients (if they exist) are replaced with their expected values conditional on the underlying. This does not affect the values of European-style options.
3. The conditional expected values from Step 2 are calculated or, more commonly, approximated. Methods based on, or related to, Gaussian approximations are often used for this step.
4. Parameter averaging techniques are used to relate the time-dependent coefficients of the SDE obtained in Step 3 to time-independent ones. This, typically, allows for a quick and direct calculation of European-style option values.

Let’s go through these 4 steps and establish the link between Gyöngy’s result and local volatility. Suppose $k = 1$ in equation (B.49) and define $X_t = \ln(S_t/S_0)$,

$\alpha(t, \omega) = \mu - 1/2\sigma^2(t, \omega)$ and $\beta(t, \omega) = \sigma(t, \omega)$. Then equation (B.49) becomes

$$d(\ln S_t) = \left(\mu - \frac{1}{2}\sigma^2(t, \omega) \right) dt + \sigma(t, \omega)dW_t \quad (\text{B.52})$$

and by Itô's lemma

$$\frac{dS_t}{S_t} = \mu dt + \sigma(t, \omega)dW_t \quad (\text{B.53})$$

which is the stochastic differential equation for a financial asset S_t with possibly stochastic volatility. Denote by S an arbitrary realisation of S_t for some $t \geq 0$ and put $x = \ln(S/S_0)$. Using Equations (B.51) we have

$$\begin{aligned} b^2(t, x) &= b^2(t, \ln(S/S_0)) = \mathbb{E}[\sigma^2(t, \omega) | S_t = S] = \sigma_{LV}(t, S) \\ a(t, x) &= a(t, \ln(S/S_0)) = \mathbb{E}\left[\mu - \frac{1}{2}\sigma^2(t, \omega) | S_t = S\right] = \mu - \frac{1}{2}\sigma_{LV}^2(t, S). \end{aligned}$$

Replacing these into equation (B.50) with $\tilde{X}_t = \ln(\tilde{S}_t/\tilde{S}_0)$ and $\tilde{S}_0 = S_0$ and using Itô's lemma, we obtain

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = \mu dt + \sigma(t, \omega)d\tilde{W}_t \quad (\text{B.54})$$

which is a SDE of \tilde{S}_t with deterministic local volatility $\sigma_{LV}(t, S)$. Thus \tilde{S}_t in the local volatility model shown in equation (B.54) has the same marginal distribution as S_t in equation (B.53) for every t . Besides, as there is a one-to-one relationship between risk-neutral marginal probabilities and the prices of standard European options, both models in Equations (B.53) and (B.54) produce the same prices for simple calls and puts after a measure change from \mathbb{P} to the risk-neutral measure.

Using Bayes relationship, the local variance becomes

$$v^2(t, S) = \frac{\mathbb{E}(\sigma^2(t, \omega)\delta(S_t - S))}{\mathbb{E}[\delta(S_t - S)]}, \quad (\text{B.55})$$

with $\delta(\bullet)$ the Dirac delta function.

Moreover Gyöngy's Theorem therefore implies that the local volatility model of equation (2.2) is in some sense the simplest diffusion model capable of doing this, i.e. reproducing the implied volatility surface.

Alexander & Nogueira (2008) states that it would be a mistake to interpret local volatility as a complete representation of the true stochastic process driving the underlying asset price. Local volatility is merely a simplification that is practically useful for describing a price process with non-constant volatility. More precisely, although the marginal distributions are the same at the time when the local volatility is calibrated, clearly S_t and \tilde{S}_t do not follow the same dynamics, hence option prices will have different dynamics under each model, and hedge ratios can differ substantially.

Bibliography

References

- Alexander, C., & Nogueira, L. (2008). *Stochastic Local Volatility*. Technical Report DP2008-02 ICMA Centre. URL: http://carolalexander.org/publish/download/DiscussionPapers/DP2008_02.pdf.
- Andersen, L., & Andreasen, J. (2000). Jump diffusion processes: Volatility smile fitting and numerical methods for option pricing. *Review of Derivatives Research*, 4, 231–261.
- Atlan, M. (2006). Localizing volatilities. *arXiv:math/0604316 [math.PR]* <http://arxiv.org/abs/math/0604316v1>, .
- Ayache, E., Henrotte, P., Nassar, S., & Wang, X. (2004). Can anyone solve the smile problem? *Wilmott Magazine*, (pp. 78–96).
- Bennett, C., & Gil, M. (2012). *Volatility Trading: Trading Volatility, Correlation, Term Structure and Skew*. Technical Report Santander Global Banking and Markets. URL: <http://cfe.cboe.com/education/TradingVolatility.pdf> equity Derivatives.
- Black, F. (1988). The holes in black-scholes. *Risk*, 1, 30–32.
- Bouzoubaa, M., & Osserein, A. (2010). *Exotic Options and Hybrids: A Guide to Structuring, Pricing and Trading*. Wiley Finance.
- Breeden, D., & Litzenberger, R. (1978). Prices of state contingent claims implicit in options price. *The Journal of Business*, 51, 621–651.
- Brigo, D., & Mercurio, F. (2001). *Interest Rate Models: Theory and Practice*. Springer.
- Brunick, G., & Shreve, S. (2013). Mimicking and itô process by a solution of a stochastic differential equation. *The Annals of Applied Probability*, 23, 1584–1628.
- Carr, P., Fisher, T., & Ruf, J. (2013). Why are quadratic normal volatility models analytically tractable? *SIAM Journal on Financial Mathematics*, 4(1), 185–202.
- Castagna, A. (2010). *FX Options and Smile Risk*. Wiley Finance.
- Clark, I. (2011). *Foreign exchange option pricing*. Wiley Finance.
- Coleman, T., Li, Y., & Verma, A. (1999). Reconstructing the unknown local volatility function. *Journal of Computational Finance*, 2, 77–100.
- Cozzi, D. (2012). *Local Stochastic Volatility Models*. Master’s thesis Politecnico di Milano Italy. URL: https://www.politesi.polimi.it/bitstream/10589/72489/1/2012_12_Cozzi.pdf.

- Derman, E., & Kani, I. (1994). Riding on a smile. *RISK Magazine*, 1, 32–39.
- Derman, E., Kani, I., & Zou, J. (1995). The local volatility surface: Unlocking the information in index option prices. *Goldman Sachs Quantitative Strategies Research Notes*, December, 32–39.
- Duffie, D. (1996). *Dynamic Asset Pricing Theory*. Princeton University Press.
- Dumas, B., Fleming, J., & Whaley, R. (1998). Whaley, implied volatility functions: Empirical tests. *Journal of Finance*, 53.
- Dupire, B. (1993). Pricing and hedging with smiles. *Paribas Capital Markets Swaps and Options Research Team*, (pp. 1–9).
- Dupire, B. (1994). Pricing with a smile. *RISK Magazine*, 1, 18–20.
- Engelmann, B., Fengler, M., & Schwendner, P. (2009). Better than its reputation: An empirical hedging analysis of the local volatility model for barrier options. *Journal of Risk*, 12, 53–77.
- Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*. New York, John Wiley and Sons.
- Gyöngy, I. (1986). Mimicking the one-dimensional marginal distributions of processes having an Itô differential. *Probability Theory and Related Fields*, 71, 501–516.
- Haugh, M. (2010). *Beyond Black-Scholes*. Technical Report Columbia University. URL: <http://www.columbia.edu/~mh2078/LocalStochasticJumpDiffusion.pdf> iEOR E4707: Financial Engineering: Continuous Time Models.
- Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 10, 327–343.
- Hull, J. (2012). *Options, Futures, and other Derivatives*. Pearson.
- Hull, J., & White, A. (1993). The pricing of options with stochastic volatilities. *Journal of Finance*, 42, 281–300.
- Kani, I., Derman, E., & Kamal, M. (1996). Trading and hedging local volatility. *Goldman Sachs Quantitative Strategies Research Notes*, August.
- Klebaner, F. (2005). *Introduction to Stochastic Calculus With Applications, Second Edition*. London, UK: Imperial College Press.
- Kotzé, A. (2003). Black-scholes or black holes? *The South African Financial Markets Journal*, 2, 8–12.

- Kotzé, A., & Joseph, A. (2009). *Constructing a South African Index Volatility Surface from Exchange Traded Data*. Technical Report Johannesburg Stock of Exchange, Equity Derivative, Report. [https://www.jse.co.za/content/JSEEducationItems/Generating South African Volatility Surface.pdf](https://www.jse.co.za/content/JSEEducationItems/Generating%20South%20African%20Volatility%20Surface.pdf).
- Kotzé, A., Labuschagne, C., Nair, M., & Padayachi, N. (2013). Arbitrage-free implied volatility surfaces for options on single stock futures. *The North American Journal of Economics and Finance*, *26*, 380–399.
- Kotzé, A., & Oosthuizen, R. (2013). Jse exotic can-do options: determining initial margins. *The South African Financial Markets Journal*, *17*. URL: <http://www.financialmarketsjournal.co.za/17thedition/jseequity.htm>.
- Kou, S. (2002). A jump diffusion model for pricing options. *Management Science*, *48*, 1086–1101.
- Lagnado, R., & Osher, S. (1997). Reconciling differences. *RISK Magazine*, *1*, 79–83.
- Linetsky, V. (1998). The path integral approach to financial modeling and options pricing. *Computational Economics*, *11*, 129–163.
- Madan, D., & Seneta, E. (1990). The variance gamma model for share market returns. *Journal of Business*, *63*, 511–524.
- Merton, R. (1976). Option pricing when the underlying stocks are discontinuous. *Journal of Financial Economics*, *5*, 125–14.
- Pitarberg, Y. (2007). Markovian projection for volatility calibration. *Risk Magazine*, *20*, 84–89.
- Rebonato, R. (2004). *Volatility and Correlation: the Perfect Hedger and the Fox, 2nd Edition*. John Wiley & Sons.
- Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance*, *49*, 771–818.
- Wilmott, P. (1998). *Derivatives*. John Wiley and Sons.