

MPRA

Munich Personal RePEc Archive

Natural volatility and option pricing

Carey, Alexander

12 January 2008

Online at <https://mpra.ub.uni-muenchen.de/6709/>
MPRA Paper No. 6709, posted 13 Jan 2008 05:16 UTC

NATURAL VOLATILITY AND OPTION PRICING

ALEXANDER CAREY

Preliminary version

January 12, 2008

In this paper we recover the Black-Scholes and local volatility pricing engines in the presence of an unspecified, fully stochastic volatility. The input volatility functions are allowed to fluctuate randomly and to depend on time to expiration in a systematic way, bringing the underlying theory in line with industry experience and practice. More generally we show that to price a European-exercise path-(in)dependent option, it is enough to model the evolution of the variance of instantaneous returns over the natural filtration of the underlying security. We call the square root of this new process *natural volatility*. We develop the associated concept of *path-conditional forward volatility*, via which the natural volatility can be directly specified in an economically meaningful way.

KEYWORDS: natural filtration, natural volatility, stochastic volatility, local volatility, path-dependent volatility, change of measure, change of filtration, martingale valuation, Black-Scholes, path-conditional forward price, path-conditional forward volatility

JEL classification: G13

1. INTRODUCTION

The widely-used Black-Scholes and local volatility models involve rigid assumptions on the spot volatility of the underlying security. The model by Black and Scholes (1973) and Merton (1973) assumes that volatility is constant, or at most some function of time. The local volatility model by Dupire (1994), Derman and Kani (1994) and Rubinstein (1994) assumes that volatility is some function of time and concurrent security price. But other than expediency, there is no compelling motivation for these assumptions. On the theoretical side, the economic basis is weak. On the empirical side, the evidence points to a volatility that displays a more complex behavior, including mean reversion and clustering.

Practical usage of the models seems to confirm that their volatility processes are misspecified. In practice, the volatility function is implied from the market prices of vanilla options, after which the model is used to price exotic contracts and manage risk. But in contravention with the theory, the implied volatility functions fluctuate with market conditions. In addition, they show a systematic dependence on time to expiration. Specifically, when calibrated to the market prices of at-the-money vanilla options, the Black-Scholes model typically predicts that volatility is to increase over time. When calibrated to a full set of vanilla options, the local volatility model predicts that volatility, as a function of the

security price, is to flatten out with the passage of time. In both cases this is contradicted by experience.

For these reasons, these two models are considered less robust than *fully stochastic* volatility models, in which the volatility process is partly or wholly driven by additional random factors. Examples include the models by Hull and White (1988), Heston (1993) and Hagan et al. (2002). In particular, operators are reluctant to use the Black-Scholes or the local volatility model for pricing derivative contracts that depend strongly on future volatility, such as *forward-start* and *cliquet* options.

Furthermore, given this empirical behavior, one cannot genuinely treat the input volatility function as the future spot volatility function, as the models would have it. Indeed, operators tend to view implied volatility as the market's estimate of future volatility, given the information available at the time of valuation (and, in the case of the local volatility model, conditional on the future security price). In fact, this unorthodox interpretation of the implied volatility function was the original basis for the idea and practice of trading volatility, that has become a centerpiece of derivatives markets. But such an interpretation is both ill-defined, and runs counter to the theory of either model.

One response has been to formally acknowledge the stochastic nature of the implied volatility functions, and incorporate it into a new, superseding model. Dupire (1996) and Derman and Kani (1998) develop *stochastic local volatility* models based on their earlier local volatility model. Schönbucher (1999) and Ledoit, Santa-Clara and Yan (2002) develop *stochastic implied volatility* models, based on the Black-Scholes model. In both cases, the original implied volatility function is to be modeled, subject to certain no-arbitrage conditions on the associated dynamics. However, conceptually the new models have little in common with the original ones, and generally produce different valuations for derivative contracts. When a valuation agrees with that of the original model, it is typically purely by construction. For example, the stochastic local volatility model agrees with the original model for any European-exercise path-independent contract, but as a bona fide fully stochastic volatility model, it otherwise generally differs (Derman and Kani 1998; Strobl 2001). As for the stochastic implied volatility model, there is generally no one instance of a calibrated Black-Scholes model that can reproduce the correct price for much more than a single derivative contract.

In this paper we recover the Black-Scholes and local volatility pricing engines in the presence of an unspecified, fully stochastic volatility. The future spot volatility function is replaced by a *path-conditional forward volatility* function. Unlike the spot volatility function of the original models, this new function is not fixed and can change with market conditions. Furthermore, it can accommodate the expiration biases found with the original models. Finally, it is closely aligned with the intuitive interpretation of the input volatility functions that is favored in practice.

More generally we show that to price a European-exercise, path-(in)dependent option, it is enough to model the *natural volatility* of the underlying security. This central construct is defined as the square root of the variance of instantaneous returns, under the equivalent martingale measure, and conditional on the history of the security price process. It turns out that the spot volatility process can be left unspecified, and any simplifying assumptions can be imposed on the natural volatility instead. We term this valuation technique *natural martingale valuation*. The results concerning the Black-Scholes and local volatility models are a direct illustration. In fact, option prices for a third type of model can also be recovered, namely *path-dependent volatility* models such as those by Hobson and Rogers (1998) and Foschi and Pascucci (2008), again without the original spot volatility assumptions.

This general result is obtained via a “change of filtration”. In standard martingale valuation, the Girsanov theorem is used to replace the expected rate of return on the security with the riskless rate of interest—a procedure known as a *change of measure*. Here we show that a similar operation can be performed for the volatility term, with natural volatility replacing spot volatility. This mechanism is implemented via a formal result.

Now from a financial viewpoint the natural volatility of the security is a somewhat abstract concept, and it is not immediately clear how to motivate a specific functional form or

specialization. This is resolved with the definition of an associated construct, the *path-conditional forward volatility* (PCFV) of the security. The square of this quantity (in fact, a function) is shown to equal the no-arbitrage premium of a forward contract on future instantaneous variance, conditional on the realization of a specific path by the underlying security price. This result follows from a more general one, which introduces the *path-conditional forward price* of an arbitrary security, and links it to the expected value of this security conditional on the history of the underlying.

The concept of path-conditional forward volatility is a natural extension of existing constructs. Dupire (1993) introduced *forward variance* as the premium in a forward contract on future spot volatility. Another quantity is the *conditional future variance* of instantaneous returns, a function of the future spot security price, which can be recovered from the initial prices of vanilla options as shown by Dupire (1996) and Derman, Kani and Kamal (1997). Other related constructs are the *corridor variance swap* (Carr and Lewis 2004) and the *conditional variance swap* (Allen, Einchcomb and Granger 2006). The special contract via which we define path-conditional forward volatility can be viewed as a conditional variance swap in which the swap is for a unique date and a vanishingly small time interval, and is triggered by the realization of a specific price trajectory by the underlying security.

This paper is organized as follows. Section 2 contains preliminary mathematical definitions and results. In Section 3 the basic financial setting is described. Section 4 reviews standard risk-neutral valuation, defines natural volatility, and introduces natural martingale valuation. Section 5 develops the concept of path-conditional forward volatility. Section 6 concludes and suggests some directions for future research.

2. MATHEMATICAL PRELIMINARIES

This section presents some basic mathematical constructs and results. First we state the definition of the Wiener process. This process is well known, but there are competing definitions, and the one we shall be using is perhaps not the most common. We then state and prove a key result, used in Section 4 to replace the volatility parameter in a stochastic differential equation (SDE).

Consider a probability space (Ω, \mathcal{F}, P) , an interval of time $\mathcal{T} = [0, T]$ and a filtration $\mathbb{F} = (\mathcal{F}_t)$. In this paper, every filtration is assumed to satisfy the usual conditions.¹

DEFINITION 2.1. *The process W is called a Wiener process relative to the probability measure P and the filtration \mathbb{F} if:*

1. $W_0 = 0$,
2. W is continuous,
3. W is a square integrable martingale relative to \mathbb{F} ,
4. if $0 \leq t < t' \leq T$ then $E_p \left[(W_{t'} - W_t)^2 \middle| \mathcal{F}_t \right] = t' - t$.

This definition is consistent with that of Lipster and Shiryaev (2001, p. 85, Definition 1). It can be shown that the increment $W_{t'} - W_t$ is independent of \mathcal{F}_t and normally distributed with mean zero and variance $t' - t$. In the special case where \mathbb{F} is the *natural filtration* of W ,² references to the filtration can be discarded and replaced with the statement that the process has independent increments. This more basic construct, often called a *standard Brownian*

¹ The *usual conditions* are that the filtration be right-continuous, and that each sigma-algebra contain the null sets. A *null set* is a set $N \subset F \in \mathcal{F}$ with $P(F) = 0$.

² The *natural filtration* of a set of stochastic processes is the minimal filtration relative to which they are adapted.

motion, is ubiquitous in the finance literature. The terminology however can vary—see Nielsen (1999, p. 50) for a survey.

A Wiener process relative to a general filtration \mathbb{F} is de facto a standard Brownian motion, while the converse is true if and only if the increment $B_{t'} - B_t$ in the Brownian motion is independent of \mathcal{F}_t , $t \in \mathcal{T}$. Bearing this in mind, the two constructs can be used interchangeably. In applications, the filtration \mathbb{F} usually governs the passage of time (Revuz and Yor 1991, p. 40). In the common case where the only source of randomness in the system is a set of standard Brownian motions, the filtration \mathbb{F} often goes unmentioned, and is then understood to be their natural filtration.

To lighten the technicalities we introduce the following notation. Given a natural number n , $\mathcal{L}_{P, \mathbb{F}}^n$ denotes the set of adapted measurable processes α such that $\int_0^T |\alpha_u|^n du < \infty$ (P -almost surely) and $\mathcal{H}_{P, \mathbb{F}}^n$ denotes the set of adapted measurable processes β such that $E_P \int_0^T |\beta_u|^n du < \infty$. Also, “ $f\left(\int_0^t \alpha_u d\beta_u\right)$ ” denotes the process $\left\{f\left(\int_0^t \alpha_u d\beta_u\right), t \in \mathcal{T}\right\}$.

We can now state the main result of this section.

PROPOSITION 2.1. *Let $\mathbb{F} = (\mathcal{F}_t)$ and $\mathbb{G} = (\mathcal{G}_t)$ be two filtrations such that $\mathcal{G}_t \subseteq \mathcal{F}_t$, $t \in \mathcal{T}$. Let W be a Wiener process relative to the probability measure P and the filtration \mathbb{F} , and let $\gamma \in \mathcal{H}_{P, \mathbb{F}}^2$. Then*

$$Z = \int \gamma_u dW_u$$

is a Wiener process relative to P and \mathbb{G} if and only if $E_P[\gamma_t^2 | \mathcal{G}_t] = 1$, $t \in \mathcal{T}$.

A proof is given in Appendix 1. Observe here that in differential form we have $dZ_t = \gamma_t dW_t$. Thus, with a suitable choice of γ , Proposition 2.1 will make it possible to change the dispersion coefficient in an SDE as in Section 4.

3. THE GENERAL FRAMEWORK

We fix a probability space (Ω, \mathcal{F}, P) and an interval of time $\mathcal{T} = [0, T]$ over which a risky security S is traded. The market information structure is represented by the filtration $\mathbb{I} = (\mathcal{I}_t)$. That is, \mathcal{I}_t is the set of all events that can be observed to be true or false by an investor by time t . We make the standard simplifying assumption that \mathcal{I}_0 contains only the null sets and their complements (Harrison and Pliska 1981). This is interpreted to mean that at time zero, investors have no information beyond the knowledge of which future events have probability zero. Under this assumption, $E[\xi | \mathcal{I}_0] = E\xi$ for any random variable ξ and probability measure equivalent to P . To avoid any distractions, the instantaneous riskless rate of interest is a constant r .

We assume there exists an equivalent martingale measure (EMM) for this market. That is, there exists a probability measure equivalent to P , under which the discounted security price is a martingale relative to the filtration \mathbb{I} (Harrison and Pliska 1981). This can be taken to be essentially equivalent to the assumption that there are no arbitrage opportunities in this market. A precise statement of this equivalence, pioneered by Harrison and Kreps (1979), is beyond the scope of this paper, and can be found in Delbaen and Schachermayer (1998) in

full generality. In general the EMM is not unique, and in this case it is common to assume that the market chooses a specific EMM for the pricing of exogenous securities (Fouque, Papanicolaou and Sircar 2000, p. 50), an approach which amounts to selecting one among several possible pricing models (Dothan 1990, pp. 294-295). We let Q denote such an EMM.

That is, the market price at time t of a time- t' payoff ξ is given by $E_Q \left[e^{-r(t'-t)} \xi \middle| \mathcal{I}_t \right]$.

Let $\mathbb{S} = (\mathcal{S}_t)$ denote the natural filtration of S , and let X be the set of sample paths of S . Define $W^{P, \mathbb{I}}$ as a Wiener process relative to the probability measure P and the filtration \mathbb{I} . Let μ be a constant, and let $\sigma \in \mathcal{H}_{Q, \mathbb{I}}^2$ be strictly positive. We assume then that the security S satisfies the SDE

$$(3.1) \quad dS_t = \mu S_t dt + \sigma_t S_t dW_t^{P, \mathbb{I}}$$

with $S_0 > 0$.

This SDE can be interpreted as describing the evolution of the security price in a world in which the passage of time is governed by the filtration \mathbb{I} . In this world, the instantaneous return at time t is normally distributed with mean μ and variance σ_t^2 per unit time, and we shall refer to μ as the *drift* and σ as the *volatility*. The process S so defined allows for the existence of an equivalent martingale measure. For the derivation of one such measure in the case where the volatility σ is an Itô process, see Fouque, Papanicolaou and Sircar (2000, pp. 46-47).

3.1. Some Nested Models

These primitives nest those of all the models mentioned thus far.

3.1.1. Constant Volatility. The spot volatility process is given by $\sigma_t = \nu$ for some positive constant ν . This is a key assumption in the model by Black and Scholes (1973).

3.1.2. Time- and Level-Dependent Volatility. The spot volatility process satisfies $\sigma_t = \nu(t, S_t)$ for some function $\nu: \mathcal{T} \times \mathbb{R} \rightarrow (0, \infty)$. The function σ can be left unspecified, and recovered numerically from an exogenous set of vanilla option prices—an approach pioneered by Dupire (1994), Derman and Kani (1994) and Rubinstein (1994). It is known as the (*deterministic*) *local volatility* model.

3.1.3. Path-Dependent Volatility. The spot volatility process satisfies $\sigma_t = \nu(t, S)$ for some function $\nu: \mathcal{T} \times X \rightarrow (0, \infty)$. Unlike the local volatility case, this specification is not operational until a specific function has been chosen. This class of models was first explored by Hobson and Rogers (1998) and Heston and Nandi (1998). Hobson and Rogers choose to model volatility as a function of the deviation of the security price from its past levels (the *offset*). Foschi and Pascucci (2008) present a variation on this approach, and relax the somewhat awkward requirement (technically not accommodated in this paper) that the security begin trading at time $-\infty$.

REMARK 3.1. Here we make ν a function of the sample path of the security S over the entire interval of time \mathcal{T} . This is essentially for notational simplicity. Note however that since the process σ is \mathbb{I} -adapted, the random variable $\nu(t, S)$ must be measurable \mathcal{I}_t , and

so the function value $v(t, x)$ depends only on the portion of the sample path x between time zero and time t .

3.1.4. Fully Stochastic Volatility. The spot volatility process is generated by additional sources of randomness besides the Wiener process $W^{P, \mathbb{I}}$. As an example, let $Z^{P, \mathbb{I}}$ denote a second Wiener process relative to the probability measure P and the filtration \mathbb{I} , that is independent of $W^{P, \mathbb{I}}$. Suppose then that spot volatility is given by $\sigma_t = \sqrt{Y_t}$, where Y is a process for which

$$dY_t = \kappa(m - Y_t)dt + v\sqrt{Y_t}(\rho dW_t^{P, \mathbb{I}} + \sqrt{1 - \rho^2} dZ_t^{P, \mathbb{I}}),$$

for suitable constants κ , m , v and ρ . This is the Heston (1993) volatility process (Fouque, Papanicolaou and Sircar 2000, pp. 40-42).

REMARK 3.2. In the models of Subsection 1, the market information structure is typically defined to be the natural filtration of the Wiener process(es) involved, although as we mentioned in section 2, this assumption is often implicit. Note that in a single-Wiener process setting, this assumption definitively precludes the possibility that volatility be fully stochastic (recall that σ must be adapted to \mathbb{I}).

4. NATURAL MARTINGALE VALUATION

This section presents the main analytical result of this paper. The objective is to price a European-exercise, path-(in)dependent derivative contract. Using Girsanov's theorem, we first derive the standard risk-neutral representation of the security price process. This representation, ordinarily used in the pricing of derivative contracts, features future spot volatility as a key determinant. We then introduce the concept of *natural volatility*, and via a change of filtration, we derive a second representation in which natural volatility replaces spot volatility. This new representation can be used in place of the standard one, with the advantages that (a) only one random factor is involved, and (b) the spot volatility process need not be specified.

Consider then a derivative contract written on the security S , with payoff $\varphi(\tau, S)$ at exercise time $\tau \in \mathcal{T}$, where $\varphi: \mathcal{T} \times X \rightarrow \mathbb{R}$. Note that the payoff can depend on the entire sample path of the security price over the interval \mathcal{T} , so that the contract can be path-dependent. However, $\varphi(\tau, S)$ is assumed to be measurable \mathcal{I}_τ , as well as Q -integrable. Further, the option has European-exercise (τ is a constant). By the definition of the equivalent martingale measure Q , the market price of the option at time zero satisfies

$$(4.1) \quad C_0 = e^{-r\tau} E_Q[\varphi(\tau, S)].$$

The valuation of the contract on the basis of this formula is known as *martingale valuation* (Nielsen 1999, p.166). In order to implement this valuation formula however, we need a suitable representation for the security price process. The next two subsections each present such a representation.

REMARK 4.1. The use of the valuation formula (4.1) is also termed *risk-neutral valuation*. The origins of the risk-neutral paradigm are to be found in Cox and Ross (1976) in the context of a complete-market model. However this designation is less compelling when the equivalent

martingale measure is not unique, and breaks down altogether when the rate of interest (or more generally the numéraire security) is correlated with the security price under the EMM.

4.1. Standard Martingale Valuation

The standard representation of the security price process is derived as follows. Since the probability measure Q is equivalent to P , as per *Girsanov's theorem* there exists a process $\lambda \in \mathcal{H}_{P, \mathbb{I}}^2$ such that

$$W^{Q, \mathbb{I}} = W^{P, \mathbb{I}} + \int \lambda_u du$$

is a Wiener process relative to Q and the filtration \mathbb{I} . Replacing in the SDE (3.1) and using the fact that the discounted security price is a martingale under Q —implying $\lambda = (\mu - r)/\sigma$ —we obtain

$$(4.2) \quad dS_t = rS_t dt + \sigma_t S_t dW_t^{Q, \mathbb{I}}.$$

For greater detail see Appendix 2. Equation (4.1) has the unique solution

$$(4.3) \quad S = S_0 \exp \left\{ \int \left(r - \frac{1}{2} \sigma_u^2 \right) du + \int \sigma_u dW_u^{Q, \mathbb{I}} \right\}$$

(Nielsen 1999, pp. 73-74, Propositions 2.18 and 2.19).

The SDE (4.2) can be interpreted as describing the motion of the security price in a mock world in which the objective probability measure is Q and the passage of time is governed by the filtration \mathbb{I} . The instantaneous return on the security is normally distributed with mean r and variance σ_t^2 per unit time.

The representation (4.3) forms the basis for standard valuation methods of the derivative contract. The contract price at time zero is obtained by inserting (4.3) into the valuation formula (4.1). The resulting expression is evaluated analytically where possible, and numerically if necessary. We shall refer to this as *standard martingale valuation*.

REMARK 4.2. Note that regardless of the method, one first needs to model the dynamics of the spot volatility process σ under the equivalent martingale measure Q . In general, different models for σ will produce different valuations for the contract.

4.2. A New Representation

We now derive a new representation for the security price process, in which the spot volatility is no longer present. For this, we define the *natural volatility* of the security S as the positive process $\bar{\sigma}$ given by

$$\bar{\sigma}_t^2 = E_Q \left[\sigma_t^2 \mid \mathcal{S}_t \right], \quad t \in \mathcal{T}.$$

We shall assume that $\bar{\sigma}$ is measurable, whereupon it can be shown that $\bar{\sigma} \in \mathcal{H}_{Q, \mathcal{S}}^2$.

The quantity $\bar{\sigma}_t^2$ can be interpreted as the variance of instantaneous returns at time t predicted by an observer whose beliefs coincide with the probability measure Q and who knows the history of the security up to time t , but nothing more. Note that the natural volatility process so defined depends on the equivalent martingale measure Q .

REMARK 4.3. For a number of models, the spot volatility is already fully determined by the history of the security price, and therefore the natural volatility equals the spot volatility. This is the case for the models 3.1.1-3.1.3, but not for the class of fully stochastic volatility models 3.1.4.

To derive the new representation, we first assume that $\sigma/\bar{\sigma} \in \mathcal{H}_{Q,\mathbb{I}}^2$. It can be verified that $E_Q \left[(\sigma_t/\bar{\sigma}_t)^2 \middle| \mathcal{S}_t \right] = 1$, $t \in \mathcal{T}$ and so by Proposition 2.1,

$$(4.4) \quad W^{Q,\mathbb{S}} = \int \frac{\sigma_u}{\bar{\sigma}_u} dW_u^{Q,\mathbb{I}}$$

is a Wiener process relative to the probability measure Q and the filtration \mathbb{S} . In differential form, this reads $\bar{\sigma}_t dW_t^{Q,\mathbb{S}} = \sigma_t dW_t^{Q,\mathbb{I}}$. Replacing (4.4) into the SDE (4.2) yields

$$(4.5) \quad dS_t = rS_t dt + \bar{\sigma}_t S_t dW_t^{Q,\mathbb{S}},$$

which has the unique solution

$$(4.6) \quad S = S_0 \exp \left\{ \int \left(r - \frac{1}{2} \bar{\sigma}_u^2 \right) du + \int \bar{\sigma}_u dW_u^{Q,\mathbb{S}} \right\}.$$

The SDE (4.5) can be viewed as describing the motion of the security price in a mock world in which the objective probability measure is Q and the passage of time is governed by the natural filtration \mathbb{S} . In this world, the instantaneous return on the security is normally distributed with mean r and variance $\bar{\sigma}_t^2$ per unit time.

The contract price at time zero can be computed by inserting (4.6) into the valuation formula (4.1). We shall call this valuation technique *natural martingale valuation*.

4.3. Practical Implications

To fix ideas we write out the two valuation formulae

$$(4.7) \quad C_0 = e^{-r\tau} E_Q \left[\varphi \left(\tau, S_0 \exp \left\{ \int \left(r - \frac{1}{2} \sigma_u^2 \right) du + \int \sigma_u dW_u^{Q,\mathbb{I}} \right\} \right) \right],$$

$$(4.8) \quad C_0 = e^{-r\tau} E_Q \left[\varphi \left(\tau, S_0 \exp \left\{ \int \left(r - \frac{1}{2} \bar{\sigma}_u^2 \right) du + \int \bar{\sigma}_u dW_u^{Q,\mathbb{S}} \right\} \right) \right].$$

The first formula corresponds to standard martingale valuation, and the second corresponds to natural martingale valuation, as described in Subsections 4.1 and 4.2 respectively. We recall that both representations proceed from the same set of primitives and assumptions, except for technical assumptions made in connection with the use of Proposition 2.1. Using formula (4.8) over (4.7) to value the derivative contract presents two potential benefits.

4.3.1. A More Economical Computation. A first advantage of formula (4.8) is that it involves only one exogenous stochastic process. Indeed, both the security S and the natural volatility $\bar{\sigma}$ are adapted to the natural filtration of the Wiener process $W^{Q,S}$. In the context of a Monte-Carlo simulation, for example, only one basic process needs to be simulated. For lattice-based valuation, only two dimensions (time and security price) are necessary. Thus, the second representation offers the prospect of a more economical computation of the contract price. In contrast, the standard valuation formula (4.7) involves two random factors: the Wiener process $W^{Q,I}$ and the spot volatility process σ .

Of course, this principle is only of use if the spot volatility is fully stochastic, as in the Heston (1993) model for example. In the case of the volatility processes of Subsections 3.1.1-3.1.3, which are wholly determined by the evolution of the security price, it offers no benefit. Further, it remains to be seen whether, for the fully stochastic volatility model in question, the natural volatility process can be determined.

4.3.2. Leaving Spot Volatility Unspecified. Another, more promising avenue, and the main emphasis in this paper, is to leave the spot volatility process unspecified altogether, and exogenously specify a natural volatility process instead. Indeed, any of the spot volatility specifications of Subsections 3.1.1-3.1.3 can be transferred to the natural volatility process instead. Thus, spot volatility is left undetermined, and only its expectation conditional on the security price history is modeled. Given that both $W^{Q,I}$ and $W^{Q,S}$ are standard Brownian motions relative to the probability measure Q , it is not difficult to convince oneself that, for a given input volatility function, formulae (4.7) and (4.8) must produce the same initial price for the contract. As an example, if we set natural volatility to be a constant $\bar{\sigma}_t = \nu$, then the price of a call option with exercise price K is

$$C_0 = S_0 N(d_1) - e^{-r\tau} KN(d_2),$$

where $d_1 = \left[\ln(S_0/K) + (r + \frac{1}{2}\nu^2)\tau \right] / (\nu\sqrt{\tau})$, $d_2 = d_1 - \nu\sqrt{\tau}$, and N is the normal distribution function. This is the classic Black-Scholes formula, except that spot volatility is replaced by natural volatility.

However, it is not obvious how to motivate exogenous specifications for natural volatility. Indeed, unlike spot volatility, it has no clear physical or economic interpretation. To remedy this, in the next section we introduce a new, closely-linked concept, with a clear financial interpretation, which will serve to specify natural volatility in an economically meaningful way.

5. PATH-CONDITIONAL FORWARD VOLATILITY

In this section we introduce a key construct of this paper. We first describe a special type of forward contract, for which both premium and delivery are conditioned on the price trajectory of the security S . The premium is termed *path-conditional forward price*. We then use this contract to define *path-conditional forward volatility*, which will allow us to specify the natural volatility in an economically meaningful way.

Consider then a contract entered into at time zero, by which two parties agree to exchange, at time t , a given security V and a pre-agreed premium, conditional on the realization of the partial sample path $\{x_u, 0 \leq u \leq t\}$ by the security S . That is, if the event $\{S_u = x_u, 0 \leq u \leq t\}$ occurs then the exchange takes place, otherwise it does not. We shall call the pre-agreed premium the *path-conditional forward price* of the security. Let the

function $\vec{V} : \mathcal{T} \times X \rightarrow \mathbb{R}$ assign to each pair (t, x) the path-conditional forward price of the security V for delivery at time t and with path $\{x_u, 0 \leq u \leq t\}$. We have the following key result.

PROPOSITION 5.1. *The path-conditional forward price function \vec{V} satisfies*

$$\vec{V}(t, S) = E_Q[V_t | \mathcal{S}_t], \quad t \in \mathcal{T}.$$

A proof is given in Appendix 3. This result is equivalently stated as

$$\vec{V}(t, x) = E_Q[V_t | S_u = x_u, 0 \leq u \leq t], \quad t \in \mathcal{T}, x \in X.$$

REMARK 5.1. In probability theory, the idea of a conditional bet dates back to de Finetti (1937) (Milne 1997). General derivative contracts with a contingent premium, among which *money-back* and *pay-later* options, have been around for some time (Kat 1994, 2001).

Consider now the case where $V_t = \sigma_t^2$. That is, the security V is a contingent claim which at time t pays the instantaneous return variance of the security S . We define $\vec{\sigma} : \mathcal{T} \times X \rightarrow (0, \infty)$ as $\vec{\sigma} = \sqrt{\vec{V}}$, and shall refer to it as *path-conditional forward volatility*. By proposition 5.1 we have

$$(5.1) \quad \vec{\sigma}(t, S)^2 = E_Q[\sigma_t^2 | \mathcal{S}_t], \quad t \in \mathcal{T},$$

that is,

$$\vec{\sigma}(t, x)^2 = E_Q[\sigma_t^2 | S_u = x_u, 0 \leq u \leq t], \quad t \in \mathcal{T}, x \in X.$$

REMARK 5.2. As in Subsection 3.1.3, although \vec{V} and $\vec{\sigma}$ are formally made functions of the sample path of the security S over the entire interval of time \mathcal{T} , the quantities $\vec{V}(t, x)$ and $\vec{\sigma}(t, x)$ depend only on the portion of the sample path x between time zero and time t .

REMARK 5.3. In this paper $\vec{\sigma}$ denotes alternatively the natural volatility process or the path-conditional forward volatility function. The meaning will normally be clear from the context.

REMARK 5.4. The existence of functions \vec{V} and $\vec{\sigma}$ satisfying Proposition 5.1 and equation (5.1) respectively, can be established independently of the financial definition we have given, using the *Doob-Dynkin lemma* (Rao and Swift 2006, p. 8, Proposition 3).

5.1 Relation to Natural Volatility

It follows immediately from (5.1) that the relation between the path-conditional forward volatility function $\vec{\sigma} : \mathcal{T} \times X \rightarrow (0, \infty)$ and the natural volatility process $\bar{\sigma}$ is

$$\vec{\sigma}(t, S) = \bar{\sigma}_t, \quad t \in \mathcal{T}.$$

This connection allows us to motivate the following specifications for natural volatility.

5.1.1. Constant PCFV. The path-conditional forward volatility function satisfies $\bar{\sigma}(t, x) = \nu$ for some constant ν . That is, the forward price at time zero of the instantaneous variance σ_t^2 , conditional on the security price path $\{x_u, 0 \leq u \leq t\}$, is the same for every sample path x and every date t . This is equivalent to specifying natural volatility as $\bar{\sigma}_t = \nu$ as in Subsection 3.1.1, and parallels the Black-Scholes model, as illustrated in Subsection 4.3.2.

5.1.2. Expiration- and Level-Dependent PCFV. The path-conditional forward volatility function has the form $\bar{\sigma}(t, x) = \nu(t, x_t)$ for some $\nu: \mathcal{T} \times \mathbb{R} \rightarrow (0, \infty)$. That is, the forward price at time zero of the instantaneous variance σ_t^2 , conditional on the security price path $\{x_u, 0 \leq u \leq t\}$, is the same for every sample path ending at x_t . This case parallels the time- and level-dependent spot volatility specification in Subsection 3.1.2. As in the general deterministic local volatility model, the function ν can be left unspecified, and recovered numerically from an exogenous set of vanilla option prices.

5.1.3. General PCFV. In the most general case, path-conditional forward volatility has—trivially—the form $\bar{\sigma}(t, x) = \nu(t, x)$ for some function $\nu: \mathcal{T} \times X \rightarrow (0, \infty)$. This specification produces the same prices as the general path-dependent spot volatility specification of Subsection 3.1.3.

6. SUMMARY AND EXTENSIONS

In this paper we introduced the concept of natural volatility, and derived a representation for the security price process based on this new quantity. For a general, fully stochastic volatility model, this representation in principle allows a more economical computation of the price of a European-exercise derivative contract, assuming that the natural volatility process is known. Further, the latter can be specified in such a way as to reproduce the pricing formulae of the Black-Scholes and local volatility models, with natural volatility in the role of spot volatility. The advantage over the original models is that spot volatility remains fully stochastic. We also developed the concept of path-conditional forward volatility, which was shown to be closely related to natural volatility, and through which the above specializations can be given a formal financial interpretation.

The results presented here could lend themselves to extensions to processes with jumps, time-dependent and stochastic interest rates, and other underlyings. In particular, it should be straightforward to recover the Black (1976) formula for the price of an interest rate caplet, again without the need for a deterministic spot volatility parameter.

The case of American exercise was not addressed in this paper. One difficulty that arises is that with a stochastic time to expiration, one must consider its interaction with the different Wiener processes involved. We must leave this issue for future research.

Finally, it remains to construct a dynamic version of the model we have developed, in which path-conditional forward volatility is made a function of calendar time as well as time to maturity and sample path. As it is, the model we have constructed is for pricing contracts at time zero. We anticipate that such an extension will require mathematical constructs that are seldom used in mathematical finance, such as random functions.

APPENDIX 1

Proof of Proposition 2.1. Beginning with the *if* statement, let us posit that $E[\gamma_t^2 | \mathcal{G}_t] = 1$ for every $t \in \mathcal{T}$. Note that since W is a Wiener process relative to the filtration \mathbb{F} , it is also a Wiener process relative to the filtration \mathbb{G} . To show that $Z = \int \gamma_u dW_u$ is a Wiener process relative to the probability measure P and the filtration \mathbb{G} , we verify conditions 1-4 from Definition 2.1.

1. $Z_0 = \int_0^0 \gamma_u dW_u = 0$ (Karatzas and Shreve 1991, p.139, Proposition 2.10).
2. Since $\gamma \in \mathcal{H}_{P, \mathbb{F}}^2 \subset \mathcal{L}_{P, \mathbb{F}}^2$ (Nielsen 1999, p. 41) the process Z has a continuous version (Nielsen 1999, proposition 1.33, p. 38).
3. Since $\gamma \in \mathcal{H}_{P, \mathbb{F}}^2$, the process Z is a martingale relative to the filtration \mathbb{F} (Nielsen 1999, p. 42, Proposition 1.37.1) and hence relative to the filtration \mathbb{G} . Furthermore,

$$E_P \left(\int_0^t \gamma_u dW_u \right)^2 = E_P \int_0^t \gamma_u^2 du < \infty$$

(Nielsen 1999, p. 42-43, Proposition 1.37.2), where the inequality follows from $\gamma \in \mathcal{H}_{P, \mathbb{F}}^2$. Thus, the process Z is square-integrable (Nielsen p. 20) and is therefore a square-integrable martingale relative to \mathbb{G} .

4. For $0 \leq t \leq t'$,

$$(A1.1) E_P \left[(Z_{t'} - Z_t)^2 | \mathcal{G}_t \right] = E_P \left[\left(\int_t^{t'} \gamma_u dW_u \right)^2 | \mathcal{G}_t \right] = E_P \left[\int_t^{t'} \gamma_u^2 du | \mathcal{G}_t \right] = \int_t^{t'} E_P \left[\gamma_u^2 | \mathcal{G}_t \right] du$$

(Nielsen 1999, pp. 42-43, Proposition 1.37.2). The integrand in the last integral simplifies to:

$$E_P \left[\gamma_u^2 | \mathcal{G}_t \right] = E_P \left[E_P \left[\gamma_u^2 | \mathcal{G}_u \right] | \mathcal{G}_t \right] = E_P \left[1 | \mathcal{G}_t \right] = 1,$$

where the first equality holds by the law of iterated expectations (Billingsley 1995, p. 448, Theorem 34.4) and the second follows from the assumption $E_P \left[\gamma_u^2 | \mathcal{G}_u \right] = 1$. Replacing in

$$(A1.1) \text{ yields } E_P \left[(Z_{t'} - Z_t)^2 | \mathcal{G}_t \right] = t' - t.$$

Thus, Proposition 2.1 holds in the ‘if’ direction. For the *only if* statement, let us assume that Z is a Wiener process relative to the probability measure P and the filtration \mathbb{G} . By Definition 2.1, $E_P \left[(Z_{t'} - Z_t)^2 | \mathcal{G}_t \right] = t' - t$. Combining this with (A1.1) yields

$$\int_t^{t'} E_P \left[\gamma_u^2 | \mathcal{G}_t \right] du = t' - t = \int_t^{t'} du$$

whenever $0 \leq t < t' \leq T$, from which $E_P[\gamma_t^2 | \mathcal{G}_t] = 1$, $t \in \mathcal{T}$ (Nielsen 1999, p. 60, Proposition 2.6). \square

APPENDIX 2

In this appendix we describe the change of measure of Subsection 4.1 in greater detail. From the assumption that the probability measures P and Q are equivalent (on \mathcal{F} , and therefore on each \mathcal{I}_t) the Radon-Nikodym derivative $E_P[dQ/dP | \mathcal{I}_t]$ can be represented as the exponential $\exp\{L_t - \frac{1}{2}\langle L, L \rangle_t\}$, where L is some continuous local martingale relative to P and \mathbb{I} (Revuz and Yor 1991, p. 304, proposition 1.6 and the preceding comment). Furthermore, by a specialization of Girsanov's theorem,

$$(A2.1) \quad W^{Q, \mathbb{I}} = W^{P, \mathbb{I}} - \langle W^{P, \mathbb{I}}, L \rangle$$

is a Wiener process relative to the probability measure Q and the filtration \mathbb{I} (Revuz and Yor 1991, p. 306, Proposition 1.12). If we assume that L satisfies $E_P\left[\exp\left\{\frac{1}{2}\langle L, L \rangle_T\right\}\right] < \infty$ (the *Novikov condition*) then L is a square-integrable martingale relative to P (Revuz and Yor 1991, pp. 308-309, Proposition 1.15 and the second comment on p. 309). It follows that L has the representation

$$L = -\int \lambda_u dW_u^{P, \mathbb{I}} + M$$

where $\lambda \in \mathcal{H}_{P, \mathbb{I}}^2$ and M is a square-integrable martingale relative to P , such that

$$(A2.2) \quad \langle W^{P, \mathbb{I}}, L \rangle = -\int \lambda_u du$$

(Lipster and Shiryaev 2001, p. 170, Theorem 5.4). Combining (A2.1) with (A2.2) yields, in differential form,

$$(A2.3) \quad dW_t^{P, \mathbb{I}} = -\lambda_t dt + dW_t^{Q, \mathbb{I}}.$$

Now by Itô's lemma (Nielsen 1999, p. 54, Theorem 2.2) the discounted price process $\bar{S} = e^{-r \cdot} S$ satisfies

$$(A2.4) \quad d\bar{S}_t = (\mu - r)\bar{S}_t dt + \sigma_t \bar{S}_t dW_t^{P, \mathbb{I}}.$$

Replacing (A2.3) into (A2.4) yields

$$d\bar{S}_t = (\mu - r - \lambda_t \sigma_t)\bar{S}_t dt + \sigma_t \bar{S}_t dW_t^{Q, \mathbb{I}}.$$

Since \bar{S} is a martingale relative to Q and \mathbb{I} —and hence also relative to the natural filtration of $W^{Q, \mathbb{I}}$ —we must have

$$\mu - r - \lambda_t \sigma_t = 0$$

(Nielsen 1999, p. 60, Proposition 2.7). That is, $\lambda = (\mu - r)/\sigma$, and using Itô's lemma again we obtain

$$dS_t = rS_t dt + \sigma_t S_t dW_t^{\mathbb{Q}, \mathbb{I}}.$$

APPENDIX 3

Proof of Proposition 5.1. First we state and prove a preliminary result. Let \mathcal{S}_t^0 denote the sigma-algebra generated by the stopped process $S_{\leq t} = \{S_u, 0 \leq u \leq t\}$.

LEMMA A3.1. *For every $H \in \mathcal{S}_t^0$,*

$$(\#) \quad \text{if } S_u(\omega) = S_u(\omega'), 0 \leq u \leq t \text{ then } \omega \in H \Rightarrow \omega' \in H.$$

Proof. We use a strategy described in Nielsen (1999, p. 321). First we seek a class \mathcal{C} of sets with the property # and which generates the sigma-algebra \mathcal{S}_t^0 . Then we show that the class of all subsets of Ω with the property # is a sigma-algebra on Ω . This class contains \mathcal{C} , and since it is a sigma-algebra it contains the sigma-algebra generated by \mathcal{C} . It will then follow that every set in \mathcal{S}_t^0 has the property #.

Beginning with the first item in this strategy, note that \mathcal{S}_t^0 is generated by the class \mathcal{C} of sets of the form

$$H = \{\omega \in \Omega : \{S_u(\omega), u = t_1, \dots, t_k\} \in \Gamma\}$$

where $0 \leq t_i \leq t$, $k \in \mathbb{N}$ and Γ is a Borel set of \mathbb{R}^k (Ibragimov and Rozanov 1978, pp. 2-3). Let $\omega \in H$, and suppose that $S_u(\omega) = S_u(\omega')$, $0 \leq u \leq t$ for some $\omega' \in \Omega$. Then clearly $S_u(\omega) = S_u(\omega')$ at the intermediate times $u = t_1, \dots, t_k$, so $\omega' \in H$. Thus, \mathcal{C} is indeed a class of sets with the property # and which generates the sigma-algebra \mathcal{S}_t^0 .

For the second item in the strategy, let \mathcal{G} denote the class of all subsets of Ω with the property #. It is easy to see that $H = \Omega$ has the property # since Ω contains every history, so

$$\Omega \in \mathcal{G}.$$

Next, let $H \in \mathcal{G}$ and $\omega \in H^c$, and suppose that $S_u(\omega) = S_u(\omega')$, $0 \leq u \leq t$ for some $\omega' \in \Omega$. Then $\omega' \in H^c$, because if $\omega' \in H$ then $\omega \in H$, contradicting the premise that $\omega \in H^c$. Thus,

$$H \in \mathcal{G} \Rightarrow H^c \in \mathcal{G}.$$

Finally, let $H_i \in \mathcal{G}$, $i=1,2,\dots$ and $\omega \in \bigcup_i H_i$. Suppose that $S_u(\omega) = S_u(\omega')$, $0 \leq u \leq t$, for some $\omega' \in \Omega$. Then $\omega \in H_j$ for some $j=1,2,\dots$, hence $\omega' \in H_j$ since $H_j \in \mathcal{G}$, and therefore $\omega' \in \bigcup_i H_i$. That is,

$$H_i \in \mathcal{G}, i=1,2,\dots \Rightarrow \bigcup_i H_i \in \mathcal{G}.$$

It follows that the class \mathcal{G} is a sigma-algebra. This proves Lemma A3.1. \square

We now turn to Proposition 5.1. Fixing t , let $X_{\leq t}$ denote the set of sample paths of the stopped process $S_{\leq t}$, with generic element $x_{\leq t} = \{x_u, 0 \leq u \leq t\}$. If $x, x' \in X$ are two sample paths such that $x_{\leq t} = x'_{\leq t}$, the definition of the function \vec{V} implies that $\vec{V}(t, x) = \vec{V}(t, x')$. It follows that there exists a function $\zeta: X_{\leq t} \rightarrow \mathbb{R}$ such that $\vec{V}(t, x) = \zeta(x_{\leq t})$ for each $x \in X$, and hence for which $\vec{V}(t, S) = \zeta(S_{\leq t})$. Since $S_{\leq t}$ generates \mathcal{S}_t^0 , by the Doob-Dynkin lemma (Rao and Swift 2006, p. 8, Proposition 3) the random variable $\vec{V}(t, S)$ is measurable \mathcal{S}_t^0 .

Next let $H \in \mathcal{S}_t^0$, and consider the following trading strategy. For each sample path $x_{\leq t}$ such that $S_{\leq t}(\omega) = x_{\leq t}$ for some $\omega \in H$, an investor takes a long position in a path-conditional forward contract on the security V with maturity t . Then if at time t the event H has occurred, the investor will pay the amount $\vec{V}(t, S)$ and receive the security with market price V_t . If H has not occurred then no contract will be triggered, because by Lemma A3.1 no history in H^c shares a path over the interval $[0, t]$ with a history in H . In summary, the investor has obtained the payoff $\mathbf{1}_H \{V_t - \vec{V}(t, S)\}$ at no cost, and so by the definition of the probability measure Q we must have

$$e^{-rt} E_Q \left[\mathbf{1}_H \{V_t - \vec{V}(t, S)\} \middle| \mathcal{I}_0 \right] = 0, \quad H \in \mathcal{S}_t^0,$$

implying

$$E_Q \left[\mathbf{1}_H \vec{V}(t, S) \right] = E \left[\mathbf{1}_H V_t \right], \quad H \in \mathcal{S}_t^0.$$

It follows that $\vec{V}(t, S)$ is the conditional expectation of V_t given \mathcal{S}_t^0 (Billingsley 1995, p. 445, Definition). It is also the conditional expectation of V_t given \mathcal{S}_t , since the latter sigma-algebra is simply the augmentation of \mathcal{S}_t^0 (Gikhman and Skorokhod 1974, p. 30). This proves Proposition 5.1. \square

REFERENCES

- ALLEN P., S. EINCHCOMB and N. GRANGER (2006): Conditional Variance Swaps. Product note, J.P. Morgan Securities, London.
- BILLINGSLEY, P. (1995): *Probability and Measure*, 3rd edition. New York: John Wiley.

- BLACK, F. (1976): The Pricing of Commodity Contracts, *J. Finan. Econ.* 3, 167-179.
- BLACK, F., and M. SCHOLES (1973): The Pricing of Options and Corporate Liabilities, *J. Polit. Economy* 81, 637-654.
- CARR, P., and K. LEWIS (2004): Corridor Variance Swaps, *Risk* February, 67-72.
- COX, J., and S. ROSS (1976): The Valuation of Options for Alternative Stochastic Processes, *J. Finan. Econ.* 3, 145-166.
- DELBAEN, F., and W. SCHACHERMAYER (1998): The Fundamental Theorem of Asset Pricing for Unbounded Stochastic Processes, *Math. Ann.* 312, 215-250.
- DERMAN, E., and I. KANI (1994): Riding on a Smile, *Risk* February, 32-39.
- DERMAN, E., I. KANI and M. KAMAL (1997): Trading and Hedging Local Volatility, *J. Finan. Eng.* 6, 233-268.
- DERMAN, E., and I. KANI (1998): Stochastic Implied Trees: Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility, *Int. J. Theoretical Appl. Finance* 1, 61-110.
- DOTHAN, M. (1990): *Prices in Financial Markets*. Oxford: Oxford University Press.
- DUPIRE, B. (1993): Arbitrage Pricing with Stochastic Volatility. Discussion paper, Paribas Capital Markets, London.
- DUPIRE, B. (1994): Pricing with a Smile, *Risk* January, 18-20.
- DUPIRE, B. (1996): A Unified Theory of Volatility. Discussion paper, Paribas Capital Markets, London.
- DE FINETTI, B. (1937): La Prévision: ses Lois Logiques, ses Sources Subjectives, *Ann. Institut Henri Poincaré* 7, 1-68.
- FOSCHI, P., and A. PASCUCCI (2008): Path Dependent Volatility, *Decis. Econ. Finance* 31, 1-20.
- FOUQUE, J.P., G. PAPANICOLAOU and K.R. SIRCAR (2000): *Derivatives in Financial Markets with Stochastic Volatility*. Cambridge: Cambridge University Press.
- GIKHMAN, I., and A. SKOROKHOD (1974): *The Theory of Stochastic Processes I*. Berlin: Springer-Verlag.
- HAGAN, P., D. KUMAR, A. LESNIEWSKI and D. WOODWARD (2002): Managing Smile Risk, *Wilmott* September, 84-108.
- HARRISON, J.M., and D. KREPS (1979): Martingales and Arbitrage in Multiperiod Securities Markets, *J. Econ. Theory* 20, 381-408.
- HARRISON, J.M. and S. PLISKA (1981): Martingales and Stochastic Integrals in the Theory of Continuous Trading, *Stoch. Process. Appl.* 11, 215-260.
- HESTON, S. (1993): A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *Rev. Finan. Stud.* 6, 327-343.
- HESTON, S., and S. NANDI (1998): Preference-Free Option Pricing with Path-Dependent Volatility: a Closed-Form Approach. Working paper 98-20, Federal Reserve Bank of Atlanta.
- HOBSON, D., and L.C.G. ROGERS (1998): Complete Models with Stochastic Volatility, *Math. Finance* 8, 27-48.

- HULL, J., and A. WHITE (1987): The Pricing of Options on Assets with Stochastic Volatilities, *J. Finance* 42, 281-300.
- IBRAGIMOV, I.A., and Y.A. ROZANOV (1978): *Gaussian Random Processes*. Berlin: Springer-Verlag.
- KARATZAS, I., and S.E. SHREVE (1991): *Brownian Motion and Stochastic Calculus*, 2nd edition. Berlin: Springer-Verlag.
- KAT, H. (1994): Contingent premium options, *J. Derivatives* 1, 44-55.
- KAT, H. (2001): *Structured Equity Derivatives*. Chichester: John Wiley.
- LEDOIT, O., P. SANTA-CLARA and S. YAN (2002): Relative Pricing of Options with Stochastic Volatility. Working paper, Anderson Graduate School of Management, Los Angeles.
- LIPSTER, R.S., and A.N. SHIRYAEV (2001): *Statistics of Random Processes I. General Theory*, 2nd edition. New York: Springer.
- MERTON, R. (1973): Theory of Rational Option Pricing, *Bell J. Econ. Manage. Sci.* 4, 141-183.
- MILNE, P. (1997): Bruno de Finetti and the Logic of Conditional Events, *Br. J. Philos. Sci.* 48, 195-232.
- NIELSEN, L.T. (1999): *Pricing and Hedging of Derivative Securities*. Oxford: Oxford University Press
- RAO, M.M., and R. SWIFT (2006): *Probability Theory with Applications*, 2nd edition. New York: Springer.
- REVUZ, D., and M. YOR (1991): *Continuous Martingales and Brownian Motion*. Berlin: Springer-Verlag.
- RUBINSTEIN, M. (1994): Implied Binomial Trees, *J. Finance* 49, 771-818.
- SCHÖNBUCHER, P. (1999): A Market Model for Stochastic Implied Volatility, *Phil. Trans. R. Soc. A* 357, 2071-2092.
- STROBL, K. (2001): On the Consistency of the Deterministic Local Volatility Function Model ('Implied Tree'), *Int. J. Theoretical Appl. Finance* 4, 545-565.