Smile-Consistent LIBOR Market Models: Theory, Approximations and Properties

Stefan Markus Kassberger

A Dissertation in Candidacy for the Degree of
Doctor of Philosophy

London School of Economics and Political Science
November 2004

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

To my parents

Acknowledgements

First of all, I wish to express my warmest thanks to my supervisor and academic teacher Prof. Dr. Rüdiger Kiesel for his confidence, guidance and support. I enjoyed my time at the Department of Statistics at the London School of Economics very much, and I am grateful to its members for creating such a stimulating and lively atmosphere. I am also indebted to Prof. Lane P. Hughston for supporting me during the first steps of my research, and to Dr. Angelos Dassios for his advice. Moreover, I want to thank Martin Forde and Martin Riesner for interesting discussions, and my colleagues at the University of Ulm for the supportive atmosphere. Finally, I would like to express my gratitude to my parents Heinz and Margaretha Kassberger for their unswerving support throughout many years of education, to my grandmother Wilhelmina Limbächer, and to M.

Financial support by the Engineering and Physical Sciences Research Council (EPSRC) is gratefully acknowledged.

Contents

List of Figures

List of Tables

Chapter 1

Implied volatility and smile modelling

"Suppose we use the standard deviation ... of possible future returns on a stock ... as a measure of its volatility. Is it reasonable to take volatility as constant over time? I think not." *Fischer Black (1976)*

1.1 Motivation

One of the pivotal assumptions the Black-Scholes-Merton theory (Black & Scholes [1973], Merton [1973]) builds on is that security prices follow a geometric Brownian motion with constant volatility. However, the quotation above suggests that this assumption was doubtful from the outset, and a large body of research shows that it is indeed inadequate (cf. e.g. Rubinstein [1994]). Log-returns in equity, foreign exchange and fixed-income markets are found to deviate heavily from normality, thus contradicting the constant-volatility premise, which would imply normally distributed log-returns.

The only non-observable parameter in the Black-Scholes formula is volatility. Given the market price of an option with a certain strike and maturity, we can find a value of the volatility parameter, the so-called *implied volatility,* such that the corresponding Black-Scholes price matches the market price of the option. The implied volatility can be ob-

tained by numerically inverting the Black-Scholes formula, and its uniqueness is guaranteed by the monotonicity of the Black-Scholes formula as a function of volatility. In an idealised Black-Scholes world, implied volatilities of options on a certain stock would be constant over all strikes and maturities. In reality, however, this is far from true. Real-world implied volatilities normally exhibit strong dependence both on strike level and on time to maturity, which is in stark contrast to the original Black-Scholes assumptions.

At this point, it is legitimate to ask why a parameter that stems from the inversion of an obviously 'incorrect' formula should deserve any attention at all. Lee [2002] gives a good answer:

> *"...it is helpful to regard the Black-Scholes implied volatility as a* language *in which to express an option price. Use of this language does not entail any belief that volatility is actually constant. A relevant analogy is the quotation of a discount bond price by giving its yield to maturity, which is the interest rate such that the observed bond price is recovered by the usual constant interest rate bond pricing formula. In no way does the use or study of bond yields entail a belief that interest rates are actually constant. As YTM is just an alternative way of expressing a bond price, so is implied volatility just an alternative way of expressing an option price. The language of implied volatility is, moreover, a* useful *alternative to raw prices. It gives a metric by which option prices can be compared across different strikes, maturities, underlyings, and observation times; and by which market prices can be compared to assessments of fair value. It is a standard in industry, to the extent that traders quote option prices in 'vol' points, and exchanges update implied volatility indices in real time."*

The challenge now is to specify models that are able to explain real-world implied volatility structures (IVSs for short).

1.2 Approaches to smile-modelling

Many different approaches have been proposed to better approximate real-world dynamics of asset prices and explain volatility smiles. In this section, we present a brief and casual overview of what we deem the most important ones and outline their main advantages and disadvantages.

We consider a frictionless financial market with a riskless bank account and a risky asset, and we assume that the price process (B_t) of the riskless bank account is described by (e^{rt}) (i.e. the bank account continuously accrues interest at a rate $r > 0$). We further posit that the price process of the risky asset (S_t) is described by

$$dS_t = rS_t\,dt + \sigma S_t\,dW_t$$

under a risk-neutral measure $\mathbb{Q}$. Then the price at time t of a European call option with strike K and maturity T is given by the Black-Scholes formula

$$C^{BS}(S_t, K, t, T, r, \sigma) = S_t N(d_1) - e^{-r(T-t)} K N(d_2),$$

with

$$d_1 = \frac{\ln(S_t/K) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}$$

and

$$d_2 = d_1 - \sigma\sqrt{T-t}.$$

Now the concept of implied volatility can be formalised as follows:

Definition 1.1. *Denote the market price at time t of a European call option with strike K and maturity T by $C(t,T,K)$. Then its Black-Scholes implied volatility is given by the unique positive solution $\sigma^{imp}(S_t, K, t, T, r)$ of the equation*

$$C^{BS}(S_t, K, t, T, r, \sigma^{imp}(S_t, K, t, T, r)) = C(t, T, K).$$

For $\sigma^{imp}(S_t, K, t, T, r)$, we use the shorthand notation ${}_t\sigma^{imp}(T, K)$.

The function ${}_t\sigma^{imp}(T,K)$ represents the implied volatility surface at time t.

A natural question arising at this point is whether one can specify alternative models that are able to reproduce or at least approximate real-world implied volatility patterns. Every alternative model implicitly gives rise to an implied volatility surface that can be obtained by calculating the model-specific call-prices $C(t,T,K)$ and then backing out the Black-Scholes implied volatilities. By comparing the shapes of model-implied volatility surfaces with volatility surfaces one typically encounters in the market, one has a natural criterion to assess the quality of a model.

1.2.1 Stochastic volatility models

A topic that has been subject of intensive research are stochastic volatility models. In this model class, (a function of) the volatility parameter is assumed to follow a stochastic process. More formally expressed, we assume that our asset-price process, considered under a risk neutral measure $\mathbb{Q}$, is governed by the following system of SDEs:

$$\begin{aligned} dS_t &= rS_t\,dt + \gamma(v_t)S_t\,dW_t^1, \\ dv_t &= \alpha(t,S_t,v_t)\,dt + \beta(t,S_t,v_t)\,dW_t^2, \end{aligned}$$

where (W_t^1) and (W_t^2) are two Brownian motions with correlation $\rho \in [-1,1]$. Under suitable regularity conditions, it can be shown that a unique solutions exists for the above system of stochastic differential equations. Here, we find ourselves in an incomplete market setting, since (v_t) is not assumed to be a traded asset. Probably the most popular parameterisations in the above general framework are the Hull & White [1987] model given by

$$\begin{aligned} dS_t &= rS_t\,dt + \sqrt{v_t}S_t\,dW_t^1, \\ dv_t &= \alpha_t v_t\,dt + \xi v_t\,dW_t^2, \end{aligned}$$

and the Heston [1993] model defined by

$$\begin{aligned} dS_t &= rS_t\,dt + \sqrt{v_t}S_t\,dW_t^1, \\ dv_t &= (\alpha_t - \kappa v_t)\,dt + \xi\sqrt{v_t}\,dW_t^2. \end{aligned}$$

Calibrating a stochastic volatility model to a typically steep short-dated implied volatility curve results in unrealistically high parameters for correlation ρ and volatility of volatility ξ. To compensate for this, calibration to a longer-dated smile requires choosing a large mean-reversion parameter κ. This suggests that stochastic volatility models are misspecified, and in particular that steep short-term implied volatility curves are not only due to stochastic volatility, but other factors, such as jumps in the asset price.

1.2.2 Models with jumps in the asset price

Incorporating jumps in the asset price process can add a certain degree of realism, because real-world asset price evolutions are far from continuous, in contrast to what diffusion-models postulate. Merton [1976] was the first to introduce jumps in the asset price by positing a jump-diffusion process of the form

$$S_t = S_0 \exp\left\{(r - \sigma^2/2)t + \sigma W_t\right\} \prod_{n=1}^{N_t} J_n,$$

where (N_t) is a Poisson process and the jumps J_n are lognormal, iid and independent of (N_t). As there is a continuum of possible jump sizes in this model, we face an incomplete market situation. Implied smile surfaces generated by the Merton model can fit steep smiles for shorter maturities (often encountered in reality) quite well, but they typically flatten out too quickly for longer maturities, thereby making it problematic to generate sufficient skewness to reproduce market smiles. Kou [2002] proposes a jump-diffusion model of the above form where the log-jump-sizes $\log(J_n)$ have an asymmetric double exponential distribution. Kou's model enhances the fit to empirical return data, and – unlike Merton's model – produces analytical pricing formulae for a range of exotic options.

Bates [1996] generalises Heston's model by adding a jump component to the asset price:

$$dS_t = rS_t\,dt + \sqrt{v_t}S_t\,dW_t^1 + S_t d\sum_{n=1}^{N_t}(J_n - 1),$$

$$dv_t = (\alpha_t - \kappa v_t)\,dt + \xi\sqrt{v_t}\,dW_t^2.$$

This formulation combines the advantages of stochastic-volatility and jump-diffusion models in so far as it is flexible enough to generate both steep skews on the short end and

moderate skews on the long end, which allows good fits to most real-world implied volatility surfaces.

Recently, exponential Lévy models have become quite popular (see e.g. Madan *et al.* [1998] for the Variance Gamma model and Carr *et al.* [2003] for a generalisation, Barndorff-Nielsen [1998] for the Normal Inverse Gaussian model, Prause [1999] for the Generalised Hyperbolic model, etc.). They incorporate jumps in the asset price and thereby – quite naturally – generate steep implied volatility structures at the short end. However, if calibrated to short maturity smiles or skews, volatility surfaces typically flatten out too quickly. If calibrated to moderate smiles or skews at longer maturities, the models produce implied volatility surfaces that are often too steep at the short end. This problem can be alleviated by using additive processes, that do not – in contrast to standard Lévy processes – feature stationary increments. Needless to say that Lévy models lead to incomplete markets.

1.2.3 Local volatility models

Probably the most natural extension of the original Black-Scholes framework is due to Dupire [1994] and Derman & Kani [1994]. They introduce deterministic volatility functions that can be both time- and state-dependent, which leads to risk-neutral asset-price dynamics of the form

$$dS_t = rS_t\,dt + \sigma(t, S_t)S_t\,dW_t,$$

with a local volatility function $\sigma(\cdot,\cdot) : \mathbb{R}^+ \times \mathbb{R}^+ \mapsto \mathbb{R}^+$ that is sufficiently regular. As opposed to stochastic volatility- and jump-models, local volatility models are complete. Moreover, local volatility models can be calibrated to exactly match implied volatility surfaces (provided these give rise to arbitrage-free option prices), whereas stochastic volatility or jump models usually cannot match all prices. As a consequence, the main purpose of a (calibrated) local volatility model is not the pricing of vanilla options or the identification of possible mispricings in the vanilla market; prices of standard options are regarded as inputs to the model. Rather, a calibrated local volatility model will typically be used to price exotics in line with vanilla options, i.e. local volatility models are mostly used as

relative pricing tools. The main problems associated with this model class is the determination of a suitable local volatility function, namely one that matches observed prices while being sufficiently realistic and regular, and the fact that local volatility models often predict future implied volatility smiles that are much flatter than current ones. Derman [2003] calls this *"... an uncomfortable and unrealistic forecast that contradicts the omnipresent nature of the skew."* We defer a more technical treatment of this issue and local volatility models in general to the following chapters.

1.2.4 Other approaches

Recently, so-called universal volatility models have been proposed that combine all the above features (stochastic volatility, jumps, local volatility) and allow even for a jump component in the volatility process. Of course, this class of models is the most realistic, but this is paid for by a large number of parameters, most of which cannot be directly observed and estimated, associated with the problem of unstable parameter estimates that can lead to unstable hedging strategies when the model is recalibrated. Again, these models lead to highly incomplete markets, which raises hedging issues.

1.3 Smile-modelling in the context of LIBOR market models

Not only can volatility smiles be observed in equity or foreign exchange markets, but also, as documented e.g. by Jarrow *et al.* [2003], in interest rate markets. These authors also find that smiles have become more pronounced after September 11, 2001, and that the standard LIBOR market model (LMM for short) – incapable of incorporating smiles – gives rise to large pricing errors and performs poorly after that date.

In recent years, the standard LMM has been extended in a variety of ways. For example, Andersen & Andreasen [2000] developed a constant elasticity of variance LMM, which falls into the local volatility category, Andersen & Brotherton-Ratcliffe [2001] and Rebonato [2004] introduced stochastic-volatility LMMs and Glasserman & Kou [2003] formulated a

jump-diffusion LMM. However, as Jarrow *et al.* [2003] find, the existing LMMs are not able to fully capture the volatility smiles observed in real-world markets. This observation motivates one of the main objectives of this dissertation: to develop fully smile-consistent LMMs.

1.4 Outline

The remainder of the dissertation is organised as follows. Chapter 2 is devoted to the study of local volatility models. We shall first derive the meanwhile classical results of Dupire, before extending them to forward options. Then, we will use our insights to develop an approximate analytical solution to the single smile problem, and numerically test the approximations we propose. Chapter 3 is dedicated to the development of the theory of what we term generalised extended LMMs, which generalise the class of extended LMMs introduced by Andersen & Andreasen [2000]. We shall argue that this new class can be calibrated to any discrete set of caplet-smiles (e.g. by using the methods presented in Chapter 2), and subsequently develop and test price-approximations for caplets and swaptions. Chapter 4 introduces Lévy-driven LMMs. We will give a novel derivation of the relations between the various forward measures, and derive the LIBOR dynamics under the terminal measure. Subsequently, we will propose an approximate Lévy-driven LMM by introducing certain simplifying assumptions. Then, issues concerning implementation will be discussed, and the approximate model is subject to numerical testing. Finally, we shall contrast the smile dynamics induced by generalised extended LMMs with those induced by Lévy-driven LMMs. Chapter 5 summarises our findings and concludes.

Chapter 2

Local volatility functions and Dupire's formula

Probably the most demanding task when using and implementing a local volatility model is determining the local volatility function. The principal problem is that one has just a finite set of options that serve as calibration instruments (namely those that are traded in the market and thus have an observable price), which apparently is not sufficient to uniquely determine a time- and level-dependent local volatility function $\sigma(t, S)$ that reproduces these prices. Typically, optimisation methods are applied to such under-determined (also called *ill-posed*) problems: Among the class of local volatility functions, the one that solves an optimisation problem for a specific objective function (and possibly satisfies some additional criteria) is chosen.

There is a considerable literature devoted to this issue, and different ways to tackle this problem have been proposed. Early approaches by Derman & Kani [1994] and Rubinstein [1994] suggest algorithms for constructing binomial or trinomial trees that are consistent with observed option prices, where consistency is attained by exploiting the degrees of freedom implicit in the construction of the trees. The local volatility function is then implicit in the option-price consistent trees. These methods are notorious for their instability in the presence of pronounced smiles and/or high interest rates. In these cases, the algorithms can lead to negative branching probabilities and failure to reproduce input

prices (cf. Barle & Cakici [1995] and Li [2001], who also propose enhancements). Another point of criticism is that these algorithms only recover the local volatility function at a discrete set of points (the tree-nodes), which only covers a triangular region of the whole (t, S) domain. Avellaneda *et al.* [1997] suggest a relative-entropy minimisation method that uses a subjectively specified prior local volatility function to construct a time- and level-dependent representation of $\sigma(\cdot, \cdot)$. This method is known to lead to local volatility functions with sharp peaks and troughs. Apart from the fact that such a behaviour of local volatility is not overly realistic, one is likely to encounter numerical problems and instabilities when using it for pricing purposes. Lagnado & Osher [1997] present a regularisation method to find a smooth function $\sigma(\cdot, \cdot)$ that minimises a function of the gradient of the local volatility function and the difference between theoretical prices and market prices. Shortcomings of this method, as pointed out by Jackson *et al.* [1999], are the high computational cost and the fact that it only generates a discrete representation of the local volatility function described by a relatively small array of nodes, which may be insufficient when pricing exotics. Jackson *et al.* [1999] represent $\sigma(\cdot, \cdot)$ by a space-time-spline that is determined by a numerical strategy that approximately minimises a functional of the difference between theoretical prices (that are determined by $\sigma(\cdot, \cdot)$) and known market vanilla prices over a range of strikes and maturities. An overview of further optimisation methods in this context can be found in Bouchouev & Isakov [1999].

A common feature of the above approaches is that they just assume the existence of a finite number of vanilla options that serve as calibration instruments. In his acclaimed articles (Dupire [1994] and Dupire [1997]), Dupire takes a different road (see also Derman & Kani [1998] for a more technical treatment). Under the assumptions that European calls of all strikes and maturities have observable prices and that the stock price follows a diffusion, he is able to show that the local volatility function is uniquely determined. Only recently, Klebaner (cf. Klebaner [2002] and Klebaner [2003]) extended Dupire's insights to the case when the stock price process is a continuous semimartingale.

The remainder of this chapter is organised as follows. First, we will provide a formal derivation of Dupire's results, originally stated for options on spot prices. Subsequently, we will extend Dupire's insights to options on forward prices. Then, we shall derive a formula

that explicitly links local and implied volatilities for forward options. This sets the stage for tackling the so-called single smile problem. We will derive analytical approximations for the aforementioned problem, and subject these to extensive testing. Sections on issues arising in the practical application of local volatility models and a summary of our results conclude.

2.1 Dupire's formula for spot options

Our formal setup is a stochastic basis $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{0\le t\le\tau}, \mathbb{P})$. The stochastic basis is assumed to satisfy the usual conditions, i.e. $\mathcal{F}_0$ contains all $\mathbb{P}$-null sets of $\mathcal{F}$, and $\mathbb{F}$ is right-continuous. We further assume that $(W_t)_{0\le t\le\tau}$ is a Brownian motion with respect to $\mathbb{F}$. The frictionless financial market under consideration has a finite trading horizon τ and consists of a riskless bank account B with price process $(B_t) = \left(\exp\left(\int_0^t r_s\,ds\right)\right)$, where $(r_t)_{0\le t\le\tau}$ is a deterministic interest rate process, and a risky asset S with spot price process $(S_t)_{0\le t\le\tau}$ which we assume to follow a one-factor diffusion-process of the form

$$dS_t = (r_t - q_t)S_t\,dt + \sigma(t, S_t)S_t\,dW_t \tag{2.1}$$

under the $\mathbb{P}$- equivalent martingale measure $\mathbb{Q}$. Here, $(q_t)_{0\le t\le\tau}$ is the deterministic process of the dividend payout rate, and $\sigma : [0,\tau]\times\mathbb{R}_+ \mapsto \mathbb{R}_+$ is always assumed to be sufficiently regular to guarantee the existence of a unique solution of (2.1).

Theorem 2.1 (Dupire). *Assume that the t-market prices $C(t,T,K)$ of European call options for all $(T,K)\in[t,\tau]\times\mathbb{R}_+$ are known and arbitrage-free, and that the derivatives $\partial C(t,T,K)/\partial T$ and $\partial^2 C(t,T,K)/\partial K^2$ exist for all $(T,K)\in[t,\tau]\times\mathbb{R}_+$. Further assume that the local volatility function $\sigma : [t,\tau]\times\mathbb{R}_+\mapsto\mathbb{R}_+$ defined by*

$$\sigma(T,K) = \sqrt{2\frac{\frac{\partial C(t,T,K)}{\partial T} + q_T C(t,T,K) + K(r_T - q_T)\frac{\partial C(t,T,K)}{\partial K}}{K^2\frac{\partial^2 C(t,T,K)}{\partial K^2}}} \tag{2.2}$$

is well-defined for all $(T,K)\in[t,\tau]\times\mathbb{R}_+$. Then σ is the unique market-consistent local volatility function in the sense that it reproduces the given market prices:

$$C(t,T,K) = \exp\left(-\int_t^T r_s\,ds\right)\mathbb{E}_{\mathbb{Q}}\left[(S_T - K)^+\middle|\,\mathcal{F}_t\right].$$

In the proof, we need the forward Kolmogorov equation (also called Fokker-Planck equation), compare e.g. Shreve [2004], p.291, Øksendal [2000], p.159, or Esser & Schlag [2001]:

Theorem 2.2. *The transition density $p(t,x;T,y)$ from state x at time t to state y at time T of a diffusion $(X_s)_{s\geq 0}$ defined by $dX_s = \mu(s,X_s)\,ds + \sigma(s,X_s)\,dW_s$ satisfies the forward Kolmogorov or Fokker-Planck equation*

$$\frac{\partial p(t,x;T,y)}{\partial T} + \frac{\partial\left[\mu(T,y)p(t,x;T,y)\right]}{\partial y} = \frac{1}{2}\frac{\partial^2\left[\sigma^2(T,y)p(t,x;T,y)\right]}{\partial y^2}$$

for fixed $(t,x) \in \mathbb{R}^+ \times \mathbb{R}$. The boundary condition is given by

$$p(t,x;t,y) = \delta(x-y),$$

where δ is the Dirac delta-function.

Proof of Theorem 2.1: We obtain the arbitrage-free t-price $V(t,S_t;T,K)$ of a European call option with strike K and maturity T by means of the risk-neutral valuation formula:

$$\begin{aligned} V(t,S_t;T,K) &= \exp\left(-\int_t^T r_s\,ds\right)\mathbb{E}_{\mathbb{Q}}\left[(S_T-K)^+\middle|\,\mathcal{F}_t\right] \\ &= \exp\left(-\int_t^T r_s\,ds\right)\mathbb{E}_{\mathbb{Q}}\left[(S_T-K)^+\middle|\,S_t\right] \\ &= \exp\left(-\int_t^T r_s\,ds\right)\int_K^\infty (S-K)p(t,S_t;T,S)\,dS. \end{aligned} \tag{2.3}$$

Here, we denoted the transition density of (S_t) by p and used that (S_t) is a diffusion-process and as such Markovian. $\exp\left(-\int_t^T r_s\,ds\right)p(t,x;T,y)$ is called the *state-price density*. Differentiating (2.3) twice with respect to K gives

$$p(t,S_t;T,K) = \exp\left(\int_t^T r_s\,ds\right)\frac{\partial^2 V(t,S_t;T,K)}{\partial^2 K}. \tag{2.4}$$

Formula (2.4) is originally due to Breeden & Litzenberger [1978]. By the forward Kolmogorov equation, the transition density p satisfies the PDE

$$\begin{aligned} &\frac{\partial p(t,S_t;T,K)}{\partial T} + \frac{\partial\left[(r_T-q_T)Kp(t,S_t;T,K)\right]}{\partial K} \\ &= \frac{1}{2}\frac{\partial^2\left[\sigma^2(T,K)K^2p(t,S_t;T,K)\right]}{\partial K^2} \end{aligned} \tag{2.5}$$

for $T \in (t, \tau]$, $K \geq 0$ and fixed (t, S_t). The boundary condition is given by

$$p(t, S; t, u) = \delta(S - u), \tag{2.6}$$

where δ is the Dirac delta-function.

We will now use representation (2.4) to express the three terms in the Fokker-Planck equation (2.5) through V, thereby eliminating p. Doing so for the first term gives

$$\begin{aligned}
&\frac{\partial p(t, S_t; T, K)}{\partial T} \\
&= \frac{\partial}{\partial T}\left[\exp\left(\int_t^T r_s\, ds\right)\frac{\partial^2 V(t, S_t; T, K)}{\partial^2 K}\right] \\
&= \exp\left(\int_t^T r_s\, ds\right)\frac{\partial^2}{\partial K^2}\left[\frac{\partial V(t, S_t; T, K)}{\partial T} + r_T V(t, S_t; T, K)\right].
\end{aligned} \tag{2.7}$$

The second term yields

$$\begin{aligned}
&\frac{\partial\left[(r_T - q_T) K p(t, S_t; T, K)\right]}{\partial K} \\
&= (r_T - q_T)\frac{\partial}{\partial K}\left[K \exp\left(\int_t^T r_s\, ds\right)\frac{\partial^2 V(t, S_t; T, K)}{\partial^2 K}\right] \\
&= (r_T - q_T)\exp\left(\int_t^T r_s\, ds\right)\frac{\partial}{\partial K}\left[K\frac{\partial^2 V(t, S_t; T, K)}{\partial^2 K}\right].
\end{aligned} \tag{2.8}$$

Finally, the third term takes the form

$$\begin{aligned}
&\frac{1}{2}\frac{\partial^2\left[\sigma^2(T, K) K^2 p(t, S_t; T, K)\right]}{\partial K^2} \\
&= \frac{1}{2}\frac{\partial^2}{\partial K^2}\left[\sigma^2(T, K) K^2 \exp\left(\int_t^T r_s\, ds\right)\frac{\partial^2 V(t, S_t; T, K)}{\partial^2 K}\right] \\
&= \frac{1}{2}\exp\left(\int_t^T r_s\, ds\right)\frac{\partial^2}{\partial K^2}\left[\sigma^2(T, K) K^2 \frac{\partial^2 V(t, S_t; T, K)}{\partial^2 K}\right].
\end{aligned} \tag{2.9}$$

Inserting (2.7), (2.8) and (2.9) into the Fokker-Planck equation (2.5) leads to

$$\begin{aligned}
&\exp\left(\int_t^T r_s\, ds\right)\frac{\partial^2}{\partial K^2}\left[\frac{\partial V(t, S_t; T, K)}{\partial T} + r_T V(t, S_t; T, K)\right] \\
&+ (r_T - q_T)\exp\left(\int_t^T r_s\, ds\right)\frac{\partial}{\partial K}\left[K\frac{\partial^2 V(t, S_t; T, K)}{\partial^2 K}\right] \\
&= \frac{1}{2}\exp\left(\int_t^T r_s\, ds\right)\frac{\partial^2}{\partial K^2}\left[\sigma^2(T, K) K^2 \frac{\partial^2 V(t, S_t; T, K)}{\partial^2 K}\right]
\end{aligned}$$

which reduces to

$$\frac{\partial^2}{\partial K^2}\left[\frac{\partial V(t,S_t;T,K)}{\partial T}+r_T V(t,S_t;T,K)\right]+(r_T-q_T)\frac{\partial}{\partial K}\left[K\frac{\partial^2 V(t,S_t;T,K)}{\partial^2 K}\right]$$
$$=\frac{1}{2}\frac{\partial^2}{\partial K^2}\left[\sigma^2(T,K)K^2\frac{\partial^2 V(t,S_t;T,K)}{\partial^2 K}\right]. \tag{2.10}$$

Boundary condition (2.6) can be alternatively expressed as

$$\frac{\partial^2 V(t,S;t,K)}{\partial^2 K}=\delta(S-K). \tag{2.11}$$

We now integrate (2.10) with respect to K to obtain

$$\frac{\partial}{\partial K}\left[\frac{\partial V(t,S_t;T,K)}{\partial T}+r_T V(t,S_t;T,K)\right]+(r_T-q_T)K\frac{\partial^2 V(t,S_t;T,K)}{\partial^2 K}$$
$$=\frac{1}{2}\frac{\partial}{\partial K}\left[\sigma^2(T,K)K^2\frac{\partial^2 V(t,S_t;T,K)}{\partial^2 K}\right]+A(T) \tag{2.12}$$

with an integration-constant $A(T)$. Integration of the boundary condition (2.11) produces

$$\frac{\partial V(t,S;t,K)}{\partial K}=H(S-K). \tag{2.13}$$

where H denotes the Heaviside-function. Integrating (2.12) and (2.13) again with respect to K finally yields

$$\frac{\partial V(t,S_t;T,K)}{\partial T}+q_T V(t,S_t;T,K)+(r_T-q_T)K\frac{\partial V(t,S_t;T,K)}{\partial K}$$
$$=\frac{1}{2}\sigma^2(T,K)K^2\frac{\partial^2 V(t,S_t;T,K)}{\partial^2 K}+A(T)K+B(T) \tag{2.14}$$

with an integration-constant $B(T)$ and boundary condition

$$V(t,S;t,K)=\max(S-K,0). \tag{2.15}$$

Following Dupire [1994], we assume that all terms involving V and its derivatives approach zero when K goes to infinity. Under this assumption, $A(T)$ and $B(T)$ must be zero, and we get

$$\frac{\partial V(t,S_t;T,K)}{\partial T}+q_T V(t,S_t;T,K)+(r_T-q_T)K\frac{\partial V(t,S_t;T,K)}{\partial K}$$
$$=\frac{1}{2}\sigma^2(T,K)K^2\frac{\partial^2 V(t,S_t;T,K)}{\partial^2 K}. \tag{2.16}$$

Substituting $C(t,T,K)$ for $V(t,S_t;T,K)$ in the forward-PDE (2.14) and the boundary condition (2.15) while assuming that all terms that involve C and its derivatives approach zero when K goes to infinity, we get

$$\frac{\partial C(t,T,K)}{\partial T} + q_T C(t,T,K) + (r_T - q_T)K\frac{\partial C(t,T,K)}{\partial K} = \frac{1}{2}\sigma^2(T,K)K^2\frac{\partial^2 C(t,T,K)}{\partial^2 K} \tag{2.17}$$

with boundary condition

$$C(t,t,K) = \max(S_t - K, 0).$$

□

Note that equation (2.16) is of a slightly different flavour from equation (2.17). While (2.16) relates our theoretical (model) prices V to a local volatility function σ (which – abstracting from our problem of determining a local volatility function – could also be taken as given), PDE (2.17) relates observed market prices to an unknown local volatility function σ. By solving (2.17) for $\sigma(T,K)$, we can therefore back out the market-consistent local volatility function, which is of the form (2.2).

In this context, it is important to mention that the 'real-world process' that generates the market prices $C(t,T,K)$ need not necessarily be a diffusion. It could as well be a jump-diffusion, a stochastic volatility process, a Lévy process etc., or need not even be known. As long as the conditions of the theorem are met, it is possible to reproduce the observed prices with a diffusion-model, no matter what the real-world process is. However, it is not possible to reproduce any set of arbitrage-free option prices with this approach, because, as Dupire put it, "diffusions cannot generate everything".[1]

Before proceeding, some further remarks are in order. In the course of the proof, we derived a *forward PDE* for the t-price $V(t,S_t;T,K)$ of a European call option (that also holds for a European put) in the *forward variables* K and T, which we restate (in slightly

[1] For a simple counterexample, see Dupire [1993].

different notation) for the convenience of the reader:[2]

$$V_T(t,S;T,K) + q_T V(t,S;T,K) + (r_T - q_T)KV_K(t,S;T,K)$$
$$= \frac{1}{2}\sigma^2(T,K)K^2V_{KK}(t,S;T,K). \tag{2.18}$$

Here, S and t are fixed. The corresponding boundary condition for the call is $V(t,S;t,K) = \max(S-K,0)$, and $V(t,S;t,K) = \max(K-S,0)$ for the put. Note that the derivation of (2.18) hinges on the fact that the risk-neutral density-function can be expressed as the second derivative of European call or put prices with respect to the strike, and that the boundary condition is a direct consequence thereof. This not the case for other types of options, which implies that the forward PDE holds only for European vanilla options.

In contrast, the fundamental Black-Scholes PDE for the t-price $\tilde{V}(t,S,T,K)$ for European options with strike K maturing in T is of the form

$$r_T\tilde{V}(t,S;T,K)$$
$$= \tilde{V}_t(t,S;T,K) + (r_T - q_T)S\tilde{V}_S(t,S;T,K) + \frac{1}{2}\sigma^2(t,S)S^2\tilde{V}_{SS}(t,S;T,K), \tag{2.19}$$

and has to be solved under the boundary condition $\tilde{V}(T,S;T,K) = \max(S-K,0)$ for calls and $\tilde{V}(T,S;T,K) = \max(K-S,0)$ for puts. The Black-Scholes PDE is a *backward PDE* in the so-called *backward variables* S and t, while K and T are fixed. A solution $\tilde{V}$ of (2.19) is therefore a function that maps a pair (t,S) to the price of a European call with fixed strike K and fixed maturity T. As opposed to the forward PDE, the backward PDE can be solved under arbitrary boundary conditions (i.e. for arbitrary payoffs in T) to yield the arbitrage-free price of the European contingent claim under consideration, and is therefore more flexible.[3] However, the use of the forward PDE renders the simultaneous valuation of a whole range of options with different strikes and maturities possible, which is an enormous computational advantage. So, apart from its usefulness when inferring local volatilities from option prices, the forward PDE is also a handy tool for the numerical solution of pricing problems, e.g. through finite-difference methods.

[2]Derivatives of V are denoted by subscripts.

[3]Depending on the type of European claim, it might be appropriate to drop the K in the notation of (2.19).

2.2 Dupire's formula for forward options

The T_0 *-forward price* of a non-dividend paying stock[4] at time t is

$$F(t,T_0) = e^{\int_t^{T_0} r_s\,ds} S_t.$$

Assuming that the spot price (S_t) follows the SDE

$$dS_t = r_t S_t\,dt + \sigma(t,S_t)S_t\,dW_t$$

under the martingale measure $\mathbb{Q}$, we can apply Itô's formula to get

$$\begin{aligned}
dF(t,T_0) &= -r_t e^{\int_t^{T_0} r_s\,ds} S_t\,dt + e^{\int_t^{T_0} r_s\,ds}\,dS_t \\
&= -r_t e^{\int_t^{T_0} r_s\,ds} S_t\,dt + r_t e^{\int_t^{T_0} r_s\,ds} S_t\,dt + \sigma(t,S_t) e^{\int_t^{T_0} r_s\,ds} S_t\,dW_t \\
&= \sigma(t,S_t) e^{\int_t^{T_0} r_s\,ds} S_t\,dW_t \\
&= \sigma(t,S_t) F(t,T_0)\,dW_t \\
&= \sigma\left(t, e^{-\int_t^{T_0} r_s\,ds} F(t,T_0)\right) F(t,T_0)\,dW_t \\
&= \sigma^F(t,F(t,T_0)) F(t,T_0)\,dW_t \qquad (2.20)
\end{aligned}$$

with

$$\sigma^F(t,x) = \sigma\left(t, xe^{-\int_t^{T_0} r_s\,ds}\right).$$

This shows that the T_0-forward price $(F(t,T_0))$ follows a $\mathbb{Q}$-martingale.

Since $F(T_0,T_0) = S_{T_0}$, it is obvious that the t-price of a European vanilla option with maturity T_0 on the T_0-forward price $(F(t,T_0))$ must equal the t-price of the corresponding option on (S_t). In case the volatility function is constant,[5] there exists a closed-form representation for the price of a European option on a forward price (or shorter a *forward option*), which is known as *Black's formula* (cf. Black [1976]).[6] Denoting the t-price of the forward option by C^{Black}, we have[7]

$$C^{Black}(F(t,T_0),K,t,T_0,(r_s),\sigma) = e^{-\int_t^{T_0} r_s\,ds}(F(t,T_0)N(d_1) - KN(d_2)),$$

[4]The extension to stocks paying dividends at a deterministic rate is straightforward and will not be considered here.

[5]This holds also if it is deterministic and only time-dependent.

[6]The derivation follows standard arguments and is therefore omitted.

[7]Here, we only need to consider the case where option maturity and T_0 agree.

with

$$d_1 = \frac{\log(F(t,T_0)/K) + \sigma^2/2(T_0 - t)}{\sigma\sqrt{T_0 - t}}$$

and

$$d_2 = d_1 - \sigma\sqrt{T_0 - t}.$$

Observe that in contrast to the Black-Scholes formula, the interest rates (r_s) do not enter the d_i terms.

We can use this observation to relate Black-Scholes implied volatilities observed for European options on (S_t) to Black implied volatilities for the corresponding forward options:

$$\begin{aligned}
&C(t,T_0,K)\\
&= C^{BS}(S_t,K,t,T_0,(r_s),{}_t\sigma^{imp}(T_0,K))\\
&= C^{Black}(F(t,T_0),K,t,T_0,(r_s),{}_t\sigma^{imp,F}(T_0,K))\\
&= C^F(t,T_0,K),
\end{aligned}$$

where $C^F(t,T_0,K)$ is the t-price of a European option on $(F(t,T_0))$ with maturity T_0 and strike K, ${}_t\sigma^{imp}(T_0,K)$ its Black-Scholes implied volatility, and ${}_t\sigma^{imp,F}(T_0,K)$ its Black implied volatility, both observed at time t. This leads us to conclude that given ${}_t\sigma^{imp}(T_0,K)$ (or the corresponding option prices), we can (numerically) compute ${}_t\sigma^{imp,F}(T_0,K)$ and vice versa.

We now establish Dupire's theorem for forward prices.

Theorem 2.3. *Assume that the t-market prices $C^F(t,T,K)$ of European call options on a forward price $(F(t,T_0))$ are known for all $(T,K) \in [t,T_0] \times \mathbb{R}_+$ and arbitrage-free, and that the derivatives $\partial C^F(t,T,K)/\partial T$ and $\partial^2 C^F(t,T,K)/\partial K^2$ exist for all $(T,K) \in [t,T_0] \times \mathbb{R}_+$. Further assume that the local volatility function $\sigma^F : [t,T_0] \times \mathbb{R}_+ \mapsto \mathbb{R}_+$ defined by*

$$\sigma^F(T,K) = \sqrt{2\frac{\frac{\partial C^F(t,T,K)}{\partial T} + r_T C^F(t,T,K)}{K^2 \frac{\partial^2 C^F(t,T,K)}{\partial K^2}}} \tag{2.21}$$

is well-defined for all $(T,K) \in [t,T_0] \times \mathbb{R}_+$. Then σ^F is the unique market-consistent

local volatility function in the sense that it reproduces the given market prices:

$$C^F(t,T,K) = \exp\left(-\int_t^T r_s\,ds\right)\mathbb{E}_{\mathbb{Q}}\left[(F(T,T_0)-K)^+\middle|\,\mathcal{F}_t\right].$$

Proof: The proof follows the same logic as that of Dupire's theorem for spot options and will thus only be sketched. The arbitrage-free t-price $V^F(t,F(t,T_0);T,K)$ of a European call option on $F(t,T_0)$ with strike K and maturity T is

$$\begin{aligned}
&V^F(t,F(t,T_0);T,K)\\
&= \exp\left(-\int_t^T r_s\,ds\right)\mathbb{E}_{\mathbb{Q}}\left[(F(T,T_0)-K)^+\middle|\,\mathcal{F}_t\right]\\
&= \exp\left(-\int_t^T r_s\,ds\right)\int_K^\infty (F-K)p^F(t,F(t,T_0);T,F)\,dF, \qquad (2.22)
\end{aligned}$$

where the risk-neutral transition density of $(F(t,T_0))$ is denoted by p^F. Differentiating (2.22) twice with respect to K gives

$$p^F(t,F(t,T_0);T,K) = \exp\left(\int_t^T r_s\,ds\right)\frac{\partial^2 V^F(t,F(t,T_0);T,K)}{\partial^2 K}. \qquad (2.23)$$

By (2.20) and the Fokker-Planck equation, the transition density p^F satisfies the PDE

$$\frac{\partial p^F(t,F(t,T_0);T,K)}{\partial T} = \frac{1}{2}\frac{\partial^2\left[\sigma^F(T,K)^2K^2p^F(t,F(t,T_0);T,K)\right]}{\partial K^2} \qquad (2.24)$$

for $T\in(t,T_0]$, $K\geq 0$ and fixed $(t,F(t,T_0))$. The boundary condition is given by

$$p^F(t,F;t,u) = \delta(F-u). \qquad (2.25)$$

Using representation (2.23) to express the two terms in the Fokker-Planck equation (2.24) through V^F and simplifying leads to

$$\begin{aligned}
&\frac{\partial^2}{\partial K^2}\left[\frac{\partial V^F(t,F(t,T_0);T,K)}{\partial T} + r_T V^F(t,F(t,T_0);T,K)\right]\\
&= \frac{1}{2}\frac{\partial^2}{\partial K^2}\left[\sigma^F(T,K)^2K^2\frac{\partial^2 V^F(t,F(t,T_0);T,K)}{\partial^2 K}\right]. \qquad (2.26)
\end{aligned}$$

Boundary condition (2.25) can be alternatively expressed as

$$\frac{\partial^2 V^F(t,F;t,K)}{\partial^2 K} = \delta(F-K). \qquad (2.27)$$

Integrating (2.26) and (2.27) twice with respect to K yields

$$\frac{\partial V^F(t,F(t,T_0);T,K)}{\partial T} + r_T V^F(t,F(t,T_0);T,K)$$
$$= \frac{1}{2}\sigma^F(T,K)^2 K^2 \frac{\partial^2 V^F(t,F(t,T_0);T,K)}{\partial^2 K} + A(T)K + B(T), \tag{2.28}$$

with integration-constants $A(T)$ and $B(T)$ and boundary condition

$$V^F(t,F;t,K) = \max(F-K,0). \tag{2.29}$$

Substituting $C^F(t,T,K)$ for $V^F(t,F(t,T_0);T,K)$ in the forward-PDE (2.28) and the boundary condition (2.29) while assuming that all terms that appear in (2.28) and involve C^F and its derivatives approach zero when K goes to infinity, we get

$$\frac{\partial C^F(t,T,K)}{\partial T} + r_T C^F(t,T,K) = \frac{1}{2}\sigma^F(T,K)^2 K^2 \frac{\partial^2 C^F(t,T,K)}{\partial^2 K}$$

with boundary condition

$$C^F(t,K,t) = \max(F(t,T_0)-K,0).$$

□

Formula (2.21) is the basis for the next section.

2.3 Linking implied and local volatilities

In the preceding section, we proved a formula that relates option prices to the corresponding local volatility function. However, in real-world options markets, options are mostly quoted in terms of their implied volatilities. Thus, when trying to back out the market-consistent local volatility function from market observables, it is convenient to have a formula at hand that directly links implied volatilities (rather than option prices) to local volatilities. In this section, we will establish such a formula for options on forward prices. For a corresponding result for spot options, compare Andersen & Brotherton-Ratcliffe [1998], Dempster & Richards [2000] or Gatheral [2003].

Theorem 2.4. *Let* ${}_t\sigma^{imp,F}(T,K) > 0$ *be the arbitrage-free implied volatility surface*[8] *observed at time* t *with* $(T,K) \in [t,T_0] \times \mathbb{R}_+$. *Then, for fixed* $F = F(t,T_0)$ *and fixed* t, *the market-consistent local volatility function* $\sigma^F(T,K)$ *is given by*

$$\sigma^F(T,K)^2 = 2\frac{s_1(T,K) + (T-t)s_2(T,K)}{s_3(T,K) + (T-t)s_4(T,K) - (T-t)^2 s_5(T,K)}, \tag{2.30}$$

with

$$\begin{aligned}
s_1(T,K) &= 2{}_t\sigma^{imp,F}(T,K)^4, \\
s_2(T,K) &= 4{}_t\sigma^{imp,F}(T,K)^3 {}_t\sigma_T^{imp,F}(T,K), \\
s_3(T,K) &= 4\left[{}_t\sigma^{imp,F}(T,K) + K\log(F/K){}_t\sigma_K^{imp,F}(T,K)\right]^2, \\
s_4(T,K) &= 4K{}_t\sigma^{imp,F}(T,K)^3\left[{}_t\sigma_K^{imp,F}(T,K) + K{}_t\sigma_{KK}^{imp,F}(T,K)\right], \\
s_5(T,K) &= K^2{}_t\sigma^{imp,F}(T,K)^4 {}_t\sigma_K^{imp,F}(T,K)^2,
\end{aligned}$$

where subscripts denote partial derivatives, which we assume to exist.

Proof: By equation (2.21)

$$\begin{aligned}
&\sigma^F(T,K)^2 \\
&= 2\frac{C_T^F(t,T,K) + r_T C^F(t,T,K)}{K^2 C_{KK}^F(t,T,K)} \\
&= 2\frac{\frac{\partial}{\partial T}\left[e^{-\int_t^T r_s\,ds}\left(FN(d_1) - KN(d_2)\right)\right] + r_T e^{-\int_t^T r_s\,ds}\left(FN(d_1) - KN(d_2)\right)}{K^2\frac{\partial^2}{\partial K^2}\left[e^{-\int_t^T r_s\,ds}\left(FN(d_1) - KN(d_2)\right)\right]} \\
&= 2\frac{e^{-\int_t^T r_s\,ds}\frac{\partial}{\partial T}\left[FN(d_1) - KN(d_2)\right]}{e^{-\int_t^T r_s\,ds}K^2\frac{\partial^2}{\partial K^2}\left[FN(d_1) - KN(d_2)\right]} \\
&= 2\frac{\frac{\partial}{\partial T}\left[FN(d_1) - KN(d_2)\right]}{K^2\frac{\partial^2}{\partial K^2}\left[FN(d_1) - KN(d_2)\right]},
\end{aligned}$$

with $d_i = d_i\left(F,K,t,T,{}_t\sigma^{imp,F}(T,K)\right)$. We introduce the notation

$$\tilde{C}^F(t,T,K) = \tilde{C}^{Black}\left(F,K,t,T,{}_t\sigma^{imp,F}(T,K)\right) = FN(d_1) - KN(d_2)$$

[8] By an arbitrage-free implied volatility surface, we mean one whose associated set of option-prices is free of arbitrage.

for the undiscounted option price, and we omit the arguments and the index t when no confusion may arise. After some standard algebraic manipulations, we get the following relations (as usual, subscripts denote partial derivatives):

$$\begin{aligned}
\tilde{C}_K^F &= \tilde{C}_K^{Black} + \sigma_K^{imp,F}\tilde{C}_\sigma^{Black},\\
\tilde{C}_{KK}^F &= \tilde{C}_{KK}^{Black} + 2\sigma_K^{imp,F}\tilde{C}_{\sigma K}^{Black} + \left(\sigma_K^{imp,F}\right)^2\tilde{C}_{\sigma\sigma}^{Black} + \sigma_{KK}^{imp,F}\tilde{C}_\sigma^{Black},\\
\tilde{C}_T^F &= \tilde{C}_T^{Black} + \sigma_T^{imp,F}\tilde{C}_\sigma^{Black},\\
\tilde{C}_\sigma^{Black} &= \phi(d_1)F\sqrt{T-t},\\
\tilde{C}_{\sigma\sigma}^{Black} &= \tilde{C}_\sigma^{Black}d_1\left(\frac{d_1}{\sigma^{imp,F}} - \sqrt{T-t}\right),\\
\tilde{C}_{\sigma K}^{Black} &= \tilde{C}_\sigma^{Black}\frac{d_1}{K\sqrt{T-t}\sigma^{imp,F}},\\
\tilde{C}_{KK}^{Black} &= \tilde{C}_\sigma^{Black}\frac{1}{K^2(T-t)\sigma^{imp,F}},\\
\tilde{C}_T^{Black} &= \tilde{C}_\sigma^{Black}\frac{\sigma^{imp,F}}{2(T-t)},
\end{aligned}$$

where ϕ denotes the density function of the standard normal distribution. It follows that

$$\begin{aligned}
\tilde{C}_T^F &= \tilde{C}_\sigma^{Black}\left(\frac{\sigma^{imp,F}}{2(T-t)} + \sigma_T^{imp,F}\right)\\
&= \frac{\tilde{C}_\sigma^{Black}}{4(T-t)\left(\sigma^{imp,F}\right)^3}\left(s_1(T,K) + (T-t)s_2(T,K)\right).
\end{aligned}$$

Moreover, we get

$$\begin{aligned}
&\tilde{C}_{KK}^F\\
&= \tilde{C}_\sigma^{Black}\left(\frac{1}{\sigma^{imp,F}K^2T} + \frac{2d_1\sigma_K^{imp}}{\sigma^{imp,F}K\sqrt{T-t}} + d_1\left(\frac{d_1}{\sigma^{imp,F}} - \sqrt{T-t}\right)\left(\sigma_K^{imp,F}\right)^2 + \sigma_{KK}^{imp,F}\right)\\
&= \frac{\tilde{C}_\sigma^{Black}}{4(T-t)K^2\left(\sigma^{imp,F}\right)^3}\left[4\left(\sigma^{imp,F}\right)^2 + 8K\log(F/K)\sigma^{imp,F}\sigma_K^{imp}\right.\\
&\quad + 4(T-t)K\left(\sigma^{imp,F}\right)^3\sigma_K^{imp} + 4K^2\left(\log(F/K)\right)^2\left(\sigma_K^{imp}\right)^2\\
&\quad \left.-(T-t)^2K^2\left(\sigma^{imp,F}\right)^4\left(\sigma_K^{imp}\right)^2 + 4(T-t)K^2\left(\sigma^{imp,F}\right)^3\sigma_{KK}^{imp}\right]\\
&= \frac{\tilde{C}_\sigma^{Black}}{4(T-t)K^2\left(\sigma^{imp,F}\right)^3}\left[s_3(T,K) + (T-t)s_4(T,K) - (T-t)^2s_5(T,K)\right]. \qquad (2.31)
\end{aligned}$$

Summing up,

$$\sigma^F(T,K)^2 = 2\frac{\tilde{C}_T^F(t,T,K)}{\tilde{K}^2 C_{KK}^F(t,T,K)} = 2\frac{s_1(T,K)+(T-t)s_2(T,K)}{s_3(T,K)+(T-t)s_4(T,K)-(T-t)^2 s_5(T,K)}.$$

We still have to make sure that the term on the right-hand side is non-negative. Since we assumed that the given implied volatility function is positive and arbitrage-free, we can conclude by standard no-arbitrage arguments that $\tilde{C}^F(t,T,K)$ is monotonically non-decreasing in T and strictly convex in K, which completes the proof. □

In case ${}_t\sigma^{imp,F}(T,K)$ is constant in K and T, i.e. ${}_t\sigma^{imp,F}(T,K) \equiv \sigma_0$, formula (2.30) reduces to $\sigma^F(T,K)^2 = \sigma_0^2$, which is exactly what one would expect. If ${}_t\sigma^{imp,F}(T,K)$ is purely time-dependent, i.e. ${}_t\sigma^{imp,F}(T,K) = {}_t\sigma^{imp,F}(T)$, (2.30) implies $\sigma^F(T)^2 = {}_t\sigma^{imp,F}(T)^2 + 2(T-t){}_t\sigma^{imp,F}(T){}_t\sigma_T^{imp,F}(T)$. Hence,

$$\begin{aligned}\frac{1}{T_0-t}\int_t^{T_0}\sigma^F(T)^2\,dT &= \frac{1}{T_0-t}\int_t^{T_0} {}_t\sigma^{imp,F}(T)^2 + 2(T-t){}_t\sigma^{imp,F}(T){}_t\sigma_T^{imp,F}(T)\,dT \\ &= \frac{1}{T_0-t}\int_t^{T_0}\frac{\partial}{\partial T}\left[(T-t){}_t\sigma^{imp,F}(T)^2\right]\,dT \\ &= {}_t\sigma^{imp,F}(T_0)^2,\end{aligned}$$

and we recover a well-known result.

A by-product of our above analysis is the following corollary that links implied volatilities to state-price densities and transition probabilities.

Corollary 2.1. *Assume the conditions of the above theorem hold true. Then*

(i) the market-implied state-price density function $\phi^F(t,T,K)$ at time t is given by

$$\begin{aligned}&\phi^F(t,T,K) \\ &= \frac{C_\sigma^{Black}}{4(T-t)K^2\left(\sigma^{imp,F}\right)^3}\left[s_3(T,K)+(T-t)s_4(T,K)-(T-t)^2 s_5(T,K)\right]. \quad (2.32)\end{aligned}$$

(ii) the market-implied transition density $p^F(t,F;T,K)$ *of the forward price process is given by*

$$p^F(t,F;T,K) = \frac{\tilde{C}_\sigma^{Black}}{4(T-t)K^2\left(\sigma^{imp,F}\right)^3}\left[s_3(T,K)+(T-t)s_4(T,K)-(T-t)^2 s_5(T,K)\right]. \quad (2.33)$$

Proof: With equation (2.23), we have

$$p^F(t,F;T,K) = e^{\int_t^T r_s\,ds} C^F_{KK}(t,T,K) = \tilde{C}^F_{KK}(t,T,K),$$

and

$$\phi^F(t,T,K) = C^F_{KK}(t,T,K) = e^{-\int_t^T r_s\,ds}\tilde{C}^F_{KK}(t,T,K),$$

and the results follow with (2.31). □

As already mentioned, similar results can also be derived for options on spot prices. However, we will confine ourselves to the results in the forward setting, as we will need only these in later chapters.

2.4 The (single) smile problem

A meanwhile almost classical problem in local volatility modelling is the (single) smile problem that consists of determining a *time-homogeneous* local volatility function $\sigma : \mathbb{R}_+ \mapsto \mathbb{R}_+$ such that the price-dynamics given by

$$dS_t = r_t S_t\,dt + \sigma(S_t)S_t\,dW_t$$

under the risk-neutral measure $\mathbb{Q}$ give rise to option prices which are consistent with the option prices observed in the market for a given and fixed maturity T_0. More precisely, given the market prices $C(t,T_0,K)$ for all $K>0$ and fixed t and T_0, we are looking for a time-homogeneous local volatility function σ such that

$$C(t,T_0,K) = \exp\left(-\int_t^{T_0} r_s\,ds\right)\mathbb{E}_{\mathbb{Q}}\left[(S_{T_0}-K)^+|\mathcal{F}_t\right]$$

holds for all $K > 0$. Alternatively, the single smile problem can be formulated for forward options: Given $C^F(t, T_0, K)$ for all $K > 0$ and fixed t and T_0, find $\sigma^F : \mathbb{R}_+ \mapsto \mathbb{R}_+$ such that with

$$dF(t, T_0) = \sigma^F(F(t, T_0))F(t, T_0)\, dW_t$$

we have

$$C^F(t, T_0, K) = \exp\left(-\int_t^{T_0} r_s\, ds\right) \mathbb{E}_{\mathbb{Q}}\left[(F(t, T_0) - K)^+ | \mathcal{F}_t\right].$$

In the following, we will adopt the latter formulation, as it does not feature a drift term and therefore is more amenable to the solution methodology we are going to propose.

By now, no exact solution to the single smile problem has been found, but a number of different approximate solution methods have been suggested. Bouchouev & Isakov [1997] and Bouchouev & Isakov [1999] reduce the identification of the unknown time-homogeneous local volatility function to the solution of a non-linear Fredholm integral equation, which they simplify by dropping terms of higher order, and then propose an iterative solution of the simplified equation. There are some downsides to this approach. First, the integral equations involved are typically non-trivial, as is the iterative solution algorithm proposed, which relies heavily on numerical techniques. Second, the solution obtained after following some steps of the iterative algorithm is – loosely speaking – just an approximate solution of the simplified problem, whose quality is unknown. Because of the approximation error, the algorithm is not well-suited for long times to maturity. Third, even if implied volatilities are given in closed form, the output of the algorithm will only be numerical.

In a more recent paper, Jiang *et al.* [2003] advocate an iterative algorithm that recovers the local volatility function from market option prices in an optimal control framework. The problems associated with this algorithm are similar to those mentioned above: Their algorithm is highly non-trivial, iterative and approximate in nature, and purely numerical.

In this section, we propose a novel approach to the single smile problem. In contrast to the existing approaches, our approach is analytical: Given a closed-form implied volatility function that is reasonably well-behaved, we are able to state an (approximate) solution to the single smile problem (i.e. a time-homogeneous local volatility function) in analytical form.

Next, we will prove two propositions that are the basis for our analytical solution method, which will subsequently be worked out in detail. Then, we will test the quality of the solutions proposed in a number of different scenarios.

2.4.1 Extending smiles to surfaces

The starting point of our analysis is a given implied volatility function observed at time t, ${}_t\sigma^{imp,F}(T_0,K)$, for a single fixed maturity $T_0 \leq \tau$. We assume that it is arbitrage-free and sufficiently well-behaved, i.e. that all necessary derivatives exist.[9]

The idea underlying our approach is to extend ${}_t\sigma^{imp,F}(T_0,K)$ to a time-homogeneous implied volatility surface, i.e. we want to *define* an implied volatility function ${}_t\hat{\sigma}^{imp,F}(T,K)$ by

$$ {}_t\hat{\sigma}^{imp,F}(T,K) = {}_t\sigma^{imp,F}(T_0,K) \quad \forall K>0, \quad \forall T \in [t,T_0]. $$

It is by no means clear that such an extension is admissible if we want to assure that the option prices implied by the newly introduced function ${}_t\hat{\sigma}^{imp,F}(T,K)$ are arbitrage-free. The following proposition sheds light on this problem.

Proposition 2.1. *Let* ${}_t\sigma^{imp,F}(T_0,K) > 0$ *be arbitrage-free. Then* ${}_t\hat{\sigma}^{imp,F}(T,K)$ *as defined above gives rise to arbitrage-free option prices for all strikes* $K > 0$ *and all maturities* $T \in (t,T_0]$.

Proof: We have shown that (see (2.31))

$$\begin{aligned} \tilde{C}^F_{KK}(T) = \frac{\tilde{C}^{Black}_{\sigma}(T)}{4(T-t)K^2\left(\sigma^{imp,F}\right)^3}\Big[&4\left(\sigma^{imp,F}\right)^2 + 8K\log(F/K)\sigma^{imp,F}\sigma^{imp,F}_K \\ &+ 4(T-t)K\left(\sigma^{imp,F}\right)^3\sigma^{imp,F}_K + 4K^2\left(\log(F/K)\right)^2\left(\sigma^{imp,F}_K\right)^2 \\ &-(T-t)^2K^2\left(\sigma^{imp,F}\right)^4\left(\sigma^{imp,F}_K\right)^2 + 4(T-t)K^2\left(\sigma^{imp,F}\right)^3\sigma^{imp,F}_{KK}\Big], \end{aligned} \tag{2.34}$$

and the above expression is positive if and only if the prices given by $\tilde{C}^F(T)$ admit no arbitrage. Thus, it suffices to show that $\tilde{C}^F_{KK}(T_0) > 0$ implies $\tilde{C}^F_{KK}(T) > 0$ for all

[9] This assumption is not at all restrictive, because in real-world applications, implied volatility functions are typically obtained by interpolating and extrapolating discrete points on the implied volatility 'curve'. Interpolating functions can easily be chosen in such a way that differentiability is ensured.

$T \in (t, T_0]$. As $\tilde{C}_\sigma^{Black}(T) > 0$, we can confine our attention to the term

$$\frac{\tilde{C}_{KK}^F(T)}{\tilde{C}_\sigma^{Black}(T)}.$$

Now

$$\begin{aligned}
&\frac{\partial}{\partial T}\left[\frac{\tilde{C}_{KK}^F(T)}{\tilde{C}_\sigma^{Black}(T)}\right] \\
&= -\frac{4\left(\sigma^{imp,F} + K\log(F/K)\sigma_K^{imp,F}\right)^2 + (T-t)^2K^2\left(\sigma^{imp,F}\right)^4\left(\sigma_K^{imp,F}\right)^2}{4(T-t)^2K^2\left(\sigma_K^{imp,F}\right)^3} \leq 0,
\end{aligned}$$

which demonstrates that the term in square brackets is monotonically decreasing in T. Combining this with the positivity of $\tilde{C}_{KK}^{Black}(T_0)$ and $\tilde{C}_\sigma^{Black}(T_0)$, we conclude that

$$\tilde{C}_{KK}^F(T) = \tilde{C}_\sigma^{Black}(T)\frac{\tilde{C}_{KK}^F(T)}{\tilde{C}_\sigma^{Black}(T)} \geq \tilde{C}_\sigma^{Black}(T)\frac{\tilde{C}_{KK}^F(T_0)}{\tilde{C}_\sigma^{Black}(T_0)} > 0,$$

which proves the claim. □

As is obvious from the proof, the extension will generally not work forward in time, i.e. an implied volatility function that is arbitrage-free for a maturity T_0 is not necessarily arbitrage-free for $T > T_0$. More precisely, for any non-flat implied volatility function that is arbitrage-free for T_0, there exists $T > T_0$ such that the related state-price density becomes negative for some K and thus offers arbitrage opportunities.[10] This can easily be understood by considering representation (2.34) of the state-price density that features a negative term that is quadratic in T.

2.4.2 Local volatility functions for time-homogeneous implied volatility surfaces

We have just proved that we can extend a given, arbitrage-free smile to an arbitrage-free and time-homogeneous implied volatility surface. The thus obtained smile-surface gives rise to a local volatility function. Now, one could suspect that a time-homogeneous implied volatility surface leads to a time-homogeneous local volatility function. Unfortunately, this

[10] Gatheral [2000] derives a similar result.

is not the case as the following proposition, which is a direct consequence of Theorem 2.4, shows.

Proposition 2.2. *Let ${}_t\sigma^{imp,F}(T_0,K) > 0$ be an arbitrage-free implied volatility function for a fixed maturity T_0 and ${}_t\hat{\sigma}^{imp,F}(T,K) := {}_t\sigma^{imp,F}(T_0,K)$ be the extrapolated arbitrage-free implied volatility surface at time t with $(T,K) \in [t,T_0] \times \mathbb{R}_+$. Then, for fixed $F = F(t,T_0)$ and fixed t, the local volatility function $\sigma^F(T,K)$ that is consistent with ${}_t\hat{\sigma}^{imp,F}(T,K)$ is given by*

$$\sigma^F(T,K)^2 = \frac{2u_1(K)}{u_2(K) + (T-t)u_3(K) - (T-t)^2 u_4(K)}, \tag{2.35}$$

with

$$\begin{aligned}
u_1(K) &= 2{}_t\sigma^{imp,F}(T_0,K)^4, \\
u_2(K) &= 4\left[{}_t\sigma^{imp,F}(T_0,K) + K\log(F/K){}_t\sigma_K^{imp,F}(T_0,K)\right]^2, \\
u_3(K) &= 4K{}_t\sigma^{imp,F}(T_0,K)^3\left[{}_t\sigma_K^{imp,F}(T_0,K) + K{}_t\sigma_{KK}^{imp,F}(T_0,K)\right], \\
u_4(K) &= K^2{}_t\sigma^{imp,F}(T_0,K)^4{}_t\sigma_K^{imp,F}(T_0,K)^2,
\end{aligned}$$

where subscripts denote partial derivatives, which we assume to exist.

By means of the local volatility function given by (2.35), we are able to exactly reproduce our input implied volatility function ${}_t\sigma^{imp,F}(T_0,K)$.

Although the above proposition tells us that the local volatility function is non-homogeneous in time, it does not give us an idea of the degree of inhomogeneity. One could suspect from the formula that inhomogeneity – especially for small to medium times to maturity $T-t$ – is not too pronounced, since for smiles that are not too pathological (after all we assume that all smiles we deal with are arbitrage-free), ${}_t\sigma_K^{imp,F}$ and ${}_t\sigma_{KK}^{imp,F}$ should be small compared to ${}_t\sigma^{imp,F}$ itself. The reader will get a feel for the behaviour of the local volatility function in the following.

2.4.3 Time-homogeneous local volatility functions

In this section, we set out to propose a solution to the single smile problem. In order to do so, we have to find a *time-homogeneous* local volatility function that reproduces (or at least

approximates) an initial, given smile ${}_t\sigma^{imp,F}(T_0, K)$ observed at time t. Formula (2.35) provides an excellent starting point for this task: If we manage to approximate $\sigma^F(T, K)$ by a time-homogeneous function $\sigma^F_{prox}(K)$ in a way that the difference between these two functions is small, we can – loosely speaking – also expect the difference between ${}_t\sigma^{imp,F}(T_0, K)$ and the implied volatilities generated by $\sigma^F_{prox}(K)$ to be small. In the following, we propose a couple of time-homogeneous approximations that differ in the degree of sophistication and approximation error. In order to make the following methods more intuitive, we apply them to a concrete example, which will accompany us throughout this section. So let $t = 0$, $T_0 = 1$ and the current T_0-forward price of an asset be $F = F(t, T_0) = 100$, and suppose that we are given an implied volatility smile for time T_0 of the form

$$ {}_t\sigma^{imp,F}(T_0, K) = 0.5 - 0.1 \exp\left[1 - \log(K/F)^2\right], $$

where F is the current (i.e. $t = 0$) forward price of the underlying. As demonstrated above, we can extend the single smile ${}_t\sigma^{imp,F}(T_0, K)$ to a surface ${}_t\sigma^{imp,F}(T, K)$ in an arbitrage-free manner by setting

$$ {}_t\sigma^{imp,F}(T, K) = {}_t\sigma^{imp,F}(T_0, K). $$

Figure 2.1 displays the such defined time-homogeneous implied volatility surface. As implied volatility surfaces are intrinsically linked with state-price density surfaces, it is appealing to visualise the state-price density surface that corresponds to Figure 2.1. This is what we do in Figure 2.2.[11] To complete the picture, we also consider the (time-inhomogeneous) local volatility function $\sigma^F(T, K)$ that corresponds to (2.1), which is given by formula (2.35) and displayed in Figure 2.3.

Comparing implied and local volatility surfaces, we observe that there are differences, both in shape and level. The most striking difference is that although the implied volatility surface is monotonically decreasing for small K, this is not the case for the corresponding local volatility surface that increases for small strike levels K before reaching a maximum. Another difference – although not very pronounced – is the inhomogeneity in time

[11] Please note that the state-price density function was cut off at 0.05.

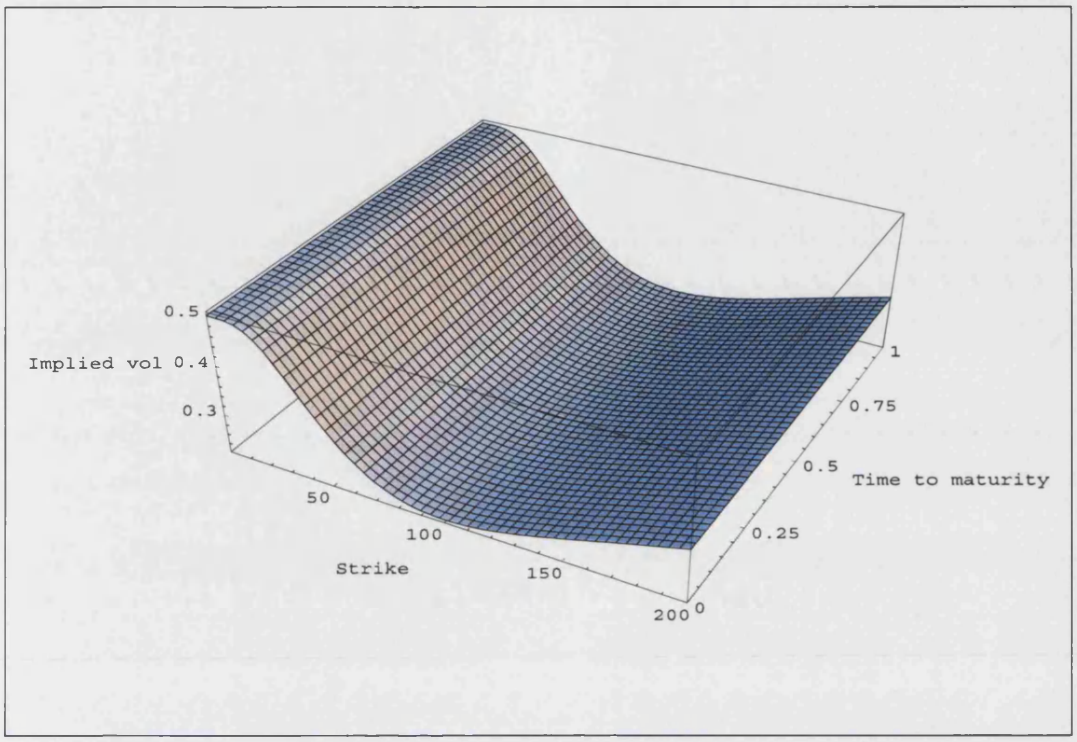

Figure 2.1: Time-homogeneous implied volatility surface

exhibited by the local volatility surface. This example indicates that the relations between implied and local volatility surfaces (even for the simple case of a time-homogeneous implied volatility surface which we are considering here) are not straightforward at all.[12]

To get a better impression of the degree time-inhomogeneity, we display the local volatility function $\sigma^F(T,K)$ for $T = 0$ and $T = 1$ in Figure 2.4. While we can observe a considerable divergence between the two curves for strikes around $K = 30$, they are close to each other for strikes over $K = 50$.

The mild degree of inhomogeneity for all but very low strikes is good news, as our objective is to find an approximating time-homogeneous local volatility function, and we can expect (at least in this example) the approximation error arising from the transition from the exact, time-inhomogeneous function to an approximating, time-homogeneous local volatility function to remain small. In the following, we discuss how this transition

[12] Derman *et al.* [1995] describe some rules of thumb for the relation between local and implied volatilities that can help develop a rough intuition.

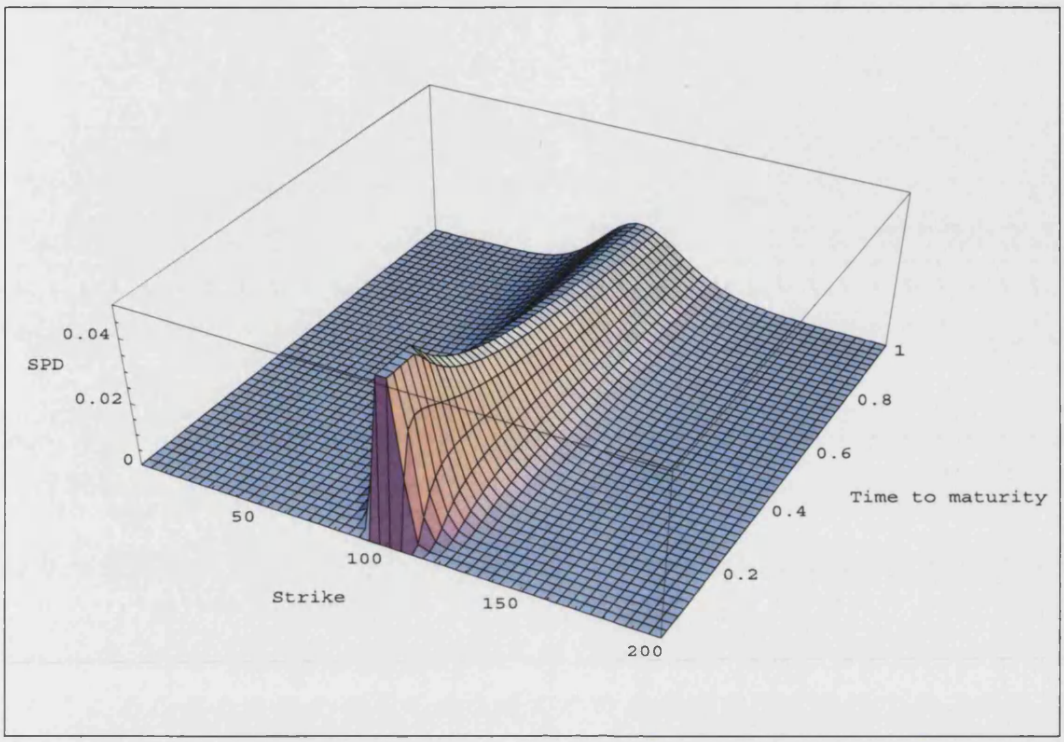

Figure 2.2: State-price density

can be accomplished.

Two simple approximations

Possibly the most obvious approximation comes about by dropping the time-inhomogeneous terms in (2.35):

$$\sigma_{p1}^F(K)^2 = \sigma^F(t,K)^2 = \frac{2u_1(K)}{u_2(K)}, \tag{P1}$$

or equivalently

$$\sigma_{p1}^F(K) = \frac{{}_t\sigma^{imp,F}(T_0,K)^2}{{}_t\sigma^{imp,F}(T_0,K) + K\log(F/K){}_t\sigma_K^{imp,F}(T_0,K)},$$

for $K > 0$. As we will see later on when we look more closely into the quality of the approximations proposed, formula (P1) – in spite of its simplicity – gives a remarkably good fit to the input data. Even a superficial look at formula (2.35) reveals the reason: For reasonably well-behaved implied volatility functions, we can expect $u_3(K)$ and $u_4(K)$ to

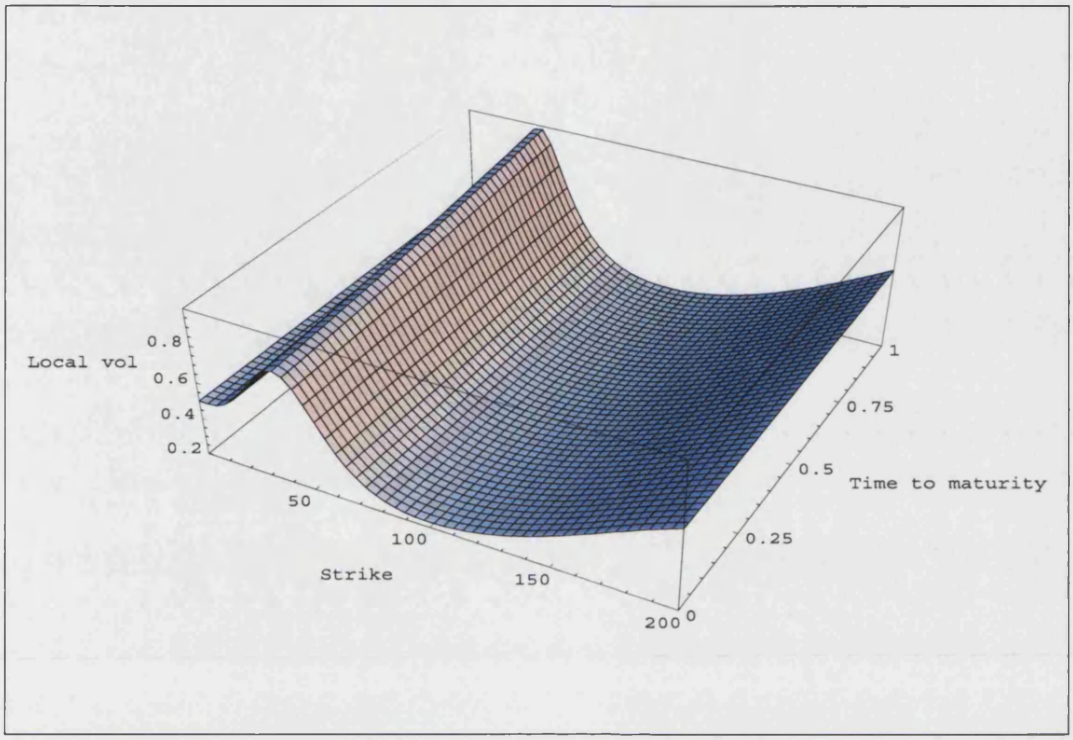

Figure 2.3: Time-inhomogeneous local volatility surface

be of a much smaller order of magnitude than $u_2(K)$, such that even for large values of $T-t$, $u_2(K)$ dominates the other two terms in the denominator. Another advantage of (P1) is that we get rid of the second derivative, which is desirable from a computational perspective.

A slightly more sophisticated approach that takes $u_3(K)$ and $u_4(K)$ into account is to evaluate (2.35) halfway between t and T_0 :

$$\begin{aligned}\sigma^F_{p2}(K) &= \sigma^F((t+T_0)/2, K) \\ &= \left(\frac{2u_1(K)}{u_2(K) + ((t+T_0)/2 - t)u_3(K) - ((t+T_0)/2 - t)^2 u_4(K)}\right)^{1/2}. \qquad \text{(P2)}\end{aligned}$$

If $\sigma^F(T,K)$ is not too far from linear in T, we can hope to get some kind of 'average' volatility along the path from t to T_0.

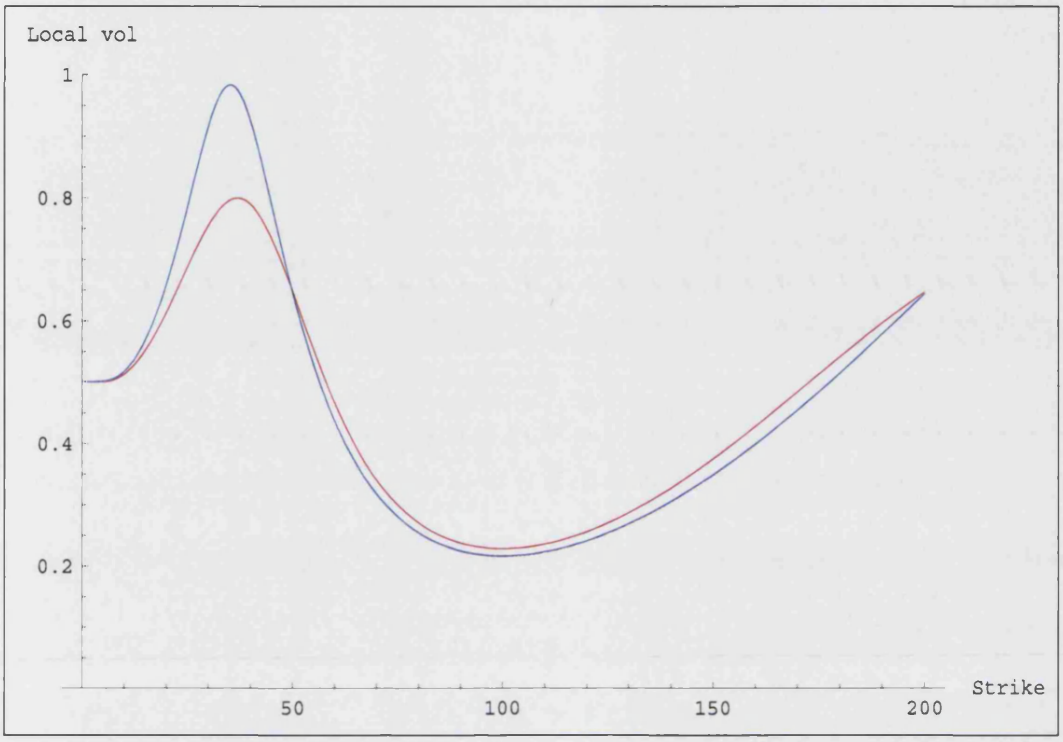

Figure 2.4: Local volatility functions $\sigma^F(0,K)$ (red line) and $\sigma^F(1,K)$ (blue line)

Integrating out T

Formula (P1) altogether neglects the information contained in $u_3(K)$ and $u_4(K)$. A natural step towards a more exact and sophisticated approximation is to incorporate u_3 and u_4 in the approximation, under the restriction that time-homogeneity is guaranteed. We can achieve this by integrating over $\sigma^F(T,K)^2$, as the following formula shows:

$$\begin{aligned}
\sigma^F_{p3}(K) &= \left(\frac{1}{T_0-t}\int_t^{T_0}\sigma^F(T,K)^2\,dT\right)^{1/2} \\
&= \left(\frac{1}{T_0-t}\int_t^{T_0}\frac{2u_1(K)}{u_2(K)+(T-t)u_3(K)-(T-t)^2u_4(K)}\,dT\right)^{1/2} \\
&= \left(\frac{2u_1(K)}{(T_0-t)\sqrt{u_3^2(K)+4u_2(K)u_4(K)}}\right. \\
&\quad \left.\cdot\left[\log\left(\frac{\sqrt{u_3^2(K)+4u_2(K)u_4(K)}-u_3(K)+2T_0u_4(K)}{\sqrt{u_3^2(K)+4u_2(K)u_4(K)}}\right)\right.\right.
\end{aligned}$$

$$- \log\left(\frac{\sqrt{u_3^2(K)+4u_2(K)u_4(K)}-u_3(K)+2tu_4(K)}{\sqrt{u_3^2(K)+4u_2(K)u_4(K)}}\right)\Bigg]\Bigg)^{1/2}$$

$$= \left(\frac{2u_1(K)}{(T_0-t)\sqrt{u_3^2(K)+4u_2(K)u_4(K)}}\right.$$

$$\left.\cdot \log\left(\left(\frac{\sqrt{u_3^2(K)+4u_2(K)u_4(K)}-u_3(K)+2T_0u_4(K)}{\sqrt{u_3^2(K)+4u_2(K)u_4(K)}-u_3(K)+2tu_4(K)}\right)\right)^{1/2}. \qquad \text{(P3)}$$

Alternatively, instead of integrating over the squared volatility, we can integrate over $\sigma^F(T,K)$:

$$\begin{aligned}
\sigma_{p4}^F(K) &= \frac{1}{T_0-t}\int_t^{T_0}\sigma^F(T,K)\,dT \\
&= \frac{1}{T_0-t}\int_t^{T_0}\sqrt{\frac{2u_1(K)}{u_2(K)+(T-t)u_3(K)-(T-t)^2u_4(K)}}\,dT \\
&= \frac{\sqrt{2u_1(K)}}{(T_0-t)\sqrt{u_4(K)}}\cdot\left[\arcsin\left(\frac{\sqrt{u_4(K)}(2T_0u_4(K)-u_3(K))}{\sqrt{4u_2(K)u_4^2(K)+u_3(K)^2u_4(K)}}\right)\right. \\
&\quad \left.-\arcsin\left(\frac{\sqrt{u_4(K)}(2tu_4(K)-u_3(K))}{\sqrt{4u_2(K)u_4^2(K)+u_3(K)^2u_4(K)}}\right)\right]. \qquad \text{(P4)}
\end{aligned}$$

Albeit a bit more complicated than (P1), above formulae are still given in closed form and can be evaluated efficiently.[13] Rather unsurprisingly, (P2) to (P4) generally lead to better approximations than formula (P1), as we will soon see.

Weighting with the state-price density

When we simply integrate T out as shown above, we neglect the fact that for a fixed strike K_0, different points on the term-structure of local volatility defined by $\sigma^F(T,K_0)$ should receive different weights. Why is that? First we observe that the state-price density function for a given maturity is tantamount to the 'discounted' marginal density of the risk-neutral forward price process.[14] As we can see in Figure 2.2 , the density is

[13] Using our previous results and standard algebra, it can be checked that all of the above expressions (in particular the integrands) are well defined.

[14] Cf. formulae (2.32) and (2.33).

concentrated around $F = 100$ for short times to maturity, and flattens out for increasing times to maturity. If we look at the SDE that governs the evolution of $(F(T,T_0))$ when using a time- and level-dependent local volatility function

$$dF(T,T_0) = \sigma^F(T,F(T,T_0))F(T,T_0)\,dW_T$$

and consider the corresponding state-price density, it becomes clear that for T close to t, the values of the local volatility function for strikes away from the money have almost no bearing on the process. For larger values of $T-t$ however, the flattening of the risk-neutral distribution implies that also local volatilities away from the money start having an impact on the process. So it is intuitively clear that the higher the state-price density in a point (T,K), the greater the impact of the value of the local volatility function in that point on $(F(T,T_0))$ and its marginal distributions. This observation lies at the heart of an approximation method which we will now make precise. We set

$$\sigma^F_{p5}(K)^2 = \int_t^{T_0} \sigma^F(T,K)^2 \frac{p^F(t,F;T,K)}{\int_t^{T_0} p^F(t,F;u,K)\,du}\,dT,$$

i.e. for fixed K, we weight every point on the term structure of local volatility with its normalised transition density, where the latter is of course based on our time-homogeneous implied volatility surface. Using previous results and simplifying yields

$$\begin{aligned}
\sigma^F_{p5}(K) &= \left(\frac{2}{K^2}\int_t^{T_0} \frac{C^F_T(t,T,K)}{C^F_{KK}(t,T,K)} \frac{C^F_{KK}(t,T,K)}{\int_t^{T_0} C^F_{KK}(t,u,K)\,du}\,dT\right)^{1/2} \\
&= \frac{\sqrt{2}}{K}\left(\frac{\int_t^{T_0} C^F_T(t,T,K)\,dT}{\int_t^{T_0} C^F_{KK}(t,u,K)du}\right)^{1/2} \\
&= \frac{\sqrt{2}}{K}\left(\frac{C^F(t,T_0,K) - C^F(t,t,K)}{\int_t^{T_0} C^F_{KK}(t,u,K)du}\right)^{1/2} \\
&= \frac{\sqrt{2}}{K}\left(\frac{C^F(t,T_0,K) - (F(t,T_0)-K)^+}{\int_t^{T_0} C^F_{KK}(t,u,K)du}\right)^{1/2}.
\end{aligned} \tag{P5}$$

Of course, we can apply the same idea to weight over $\sigma^F(T,K)$:

$$\begin{aligned}
\sigma^F_{p6}(K) &= \int_t^{T_0} \sigma^F(T,K) \frac{p^F(t,F;T,K)}{\int_t^{T_0} p^F(t,F;u,K)\,du}\,dT \\
&= \frac{\sqrt{2}}{K}\int_t^{T_0} \sqrt{\frac{C^F_T(t,T,K)}{C^F_{KK}(t,T,K)}} \frac{C^F_{KK}(t,T,K)}{\int_t^{T_0} C^F_{KK}(t,u,K)\,du}\,dT
\end{aligned}$$

$$= \frac{\sqrt{2}}{K} \frac{\int_t^{T_0} \sqrt{C_T^F(t,T,K) C_{KK}^F(t,T,K)}\, dT}{\int_t^{T_0} C_{KK}^F(t,u,K)\, du}. \tag{P6}$$

When it comes to practical applications, the drawback of the last two approximations – although given in closed form – is that they have to be evaluated by numerical integration. We will see soon whether this effort pays off in terms of goodness of fit to given smiles. Before proceeding, we remark that all the proposed approximations are able to recover the correct local volatility function for flat (i.e. state-independent) implied volatility functions $\sigma^{imp,F}$.

2.4.4 Testing the time-homogenous approximations

In this chapter, we perform extensive tests for all the proposed approximations. We investigate their quality in six different scenarios, each of which is characterised by a different combination of maturity T_0 and implied volatility function. As in our example above, we assume $t = 0$ and the current T_0-forward price of the asset to be $F = 100$. The scenarios are summarised in Table (2.1). The table shows that we consider both

T_0	skew	smile
$\frac{1}{12}$	${}_t\sigma^{imp,F}(T_0,K) = 0.2 + 2e^{-2\frac{K}{F}}$	${}_t\sigma^{imp,F}(T_0,K) = 1.2 - 0.3e^{1-\left(\log\frac{K}{F}\right)^2}$
1	${}_t\sigma^{imp,F}(T_0,K) = 0.2 + 0.4e^{-\frac{K}{F}}$	${}_t\sigma^{imp,F}(T_0,K) = 0.5 - 0.1e^{1-\left(\log\frac{K}{F}\right)^2}$
10	${}_t\sigma^{imp,F}(T_0,K) = 0.2 + 0.1e^{-0.5\frac{K}{F}}$	${}_t\sigma^{imp,F}(T_0,K) = 0.25 - 0.025e^{1-\left(\log\frac{K}{F}\right)^2}$

Table 2.1: Test scenarios

smile-shaped and skewed implied volatility functions for short, medium and long times to maturity. Thereby, we hope to cover most cases of practical interest. We intentionally choose quite extreme implied volatility patterns that feature pronounced changes in level

and non-negligible first and second derivatives. This can be regarded as stress-testing our methodology, the rationale being that a method that works for such extreme input-data should perform even better in most real-world situations, which can safely be assumed to be more moderate.

Before we come to the numerical results, we first outline our testing methodology. First, we calculate the time-homogeneous local volatility functions (P1) to (P6), which we do analytically for (P1) to (P4) and by numerical integration for (P5) and (P6). Then we use (P1) to (P6) as inputs for a Crank-Nicolson finite-difference algorithm implemented in C++, by means of which we compute the prices of standard European calls on $(F(t,T_0))$ for strike prices $10,20,\ldots,200$ for $T_0 = 1/12$ and strike prices $10,20,\ldots,300$ for $T_0 = 1$ and $T_0 = 10$. For the finite-difference grid, we use 200 time-steps for $T_0 = 1/12$, 1000 time-steps for $T_0 = 1$, and 10000 time-steps for $T_0 = 10$. For all maturities, we choose 10000 spatial steps, resulting in finite-difference grids with $2 \cdot 10^6$, 10^7 and 10^8 grid-points, respectively. This might seem excessive at first, but the objective of our tests is to be able to assess the approximation error that stems from the substitution of the exact local volatility function by our approximations (P1) to (P6). Therefore, we have to make sure that the discretization error arising from the use of a finite-difference method is kept as low as possible, which we achieve by choosing very fine grids. We also conducted tests with even finer grids that in most cases lead to slightly lower overall errors (by overall error we mean the sum of approximation- and discretization error). This indicates that the errors reported in the following are conservative.[15] We then use a Newton-Raphson-method to numerically invert the option prices computed with our algorithm (henceforth termed 'model prices') to back out the corresponding implied volatilities ('model implied volatilities') for a range of strikes that depends on the time to maturity of the option under consideration.[16] Thus, for all six scenarios, we obtain the model prices and model

[15]Even if we postulate that finer grids lead to lower discretization errors, it is not imperative that they also reduce the overall error, which results both from the use of an approximate local volatility function and the discretization error due to the finite-difference method. This is because approximation error and discretization error can also offset each other.

[16]Here we have to bear in mind that the vega on an option tends to zero if we move away from the money, or, for options away from the money, if the time to maturity goes to zero. Expressed differently, the price of an option that is far in- or out-of-the-money or one that is not at-the-money and about to expire, is

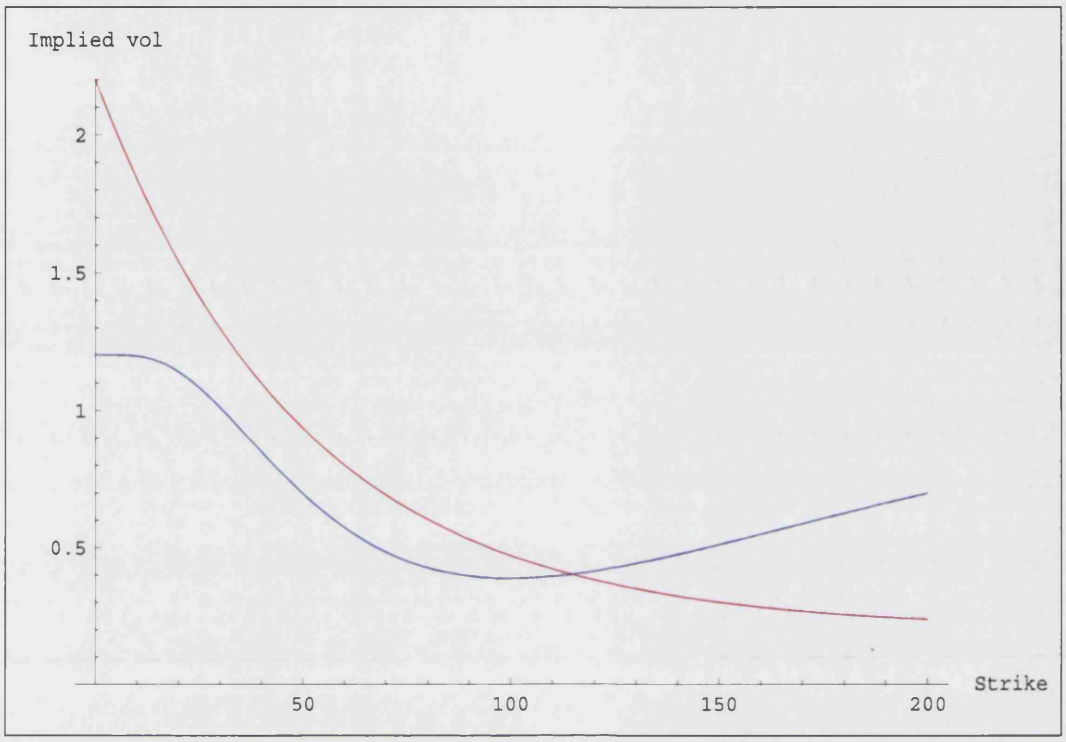

Figure 2.5: Implied volatility skew (red) and smile (blue) for $T_0 = 1/12$

implied volatilities for each of our approximation methods (P1) to (P6). Finally, we plot the model implied volatilities against the 'true' ones, and, to get a better impression of the approximation error, the differences between them. The numerical values underlying the graphs are reported in Appendix A, along with the exact and model prices (respectively their differences).

The following results are grouped according to time to maturity.

Short time to maturity: $T_0 = 1/12$

The implied volatility functions for $T_0 = 1/12$ are shown in Figure 2.5. Obviously, the levels and slopes of both implied volatility curves are quite extreme, yielding good test-cases.

insensitive to changes in (implied) volatility, so that it is not possible to reliably invert the price-formula to back out the implied volatility. In the literature, this fact is often paraphrased as the information-content of implied volatilities of options far away from the money being low.

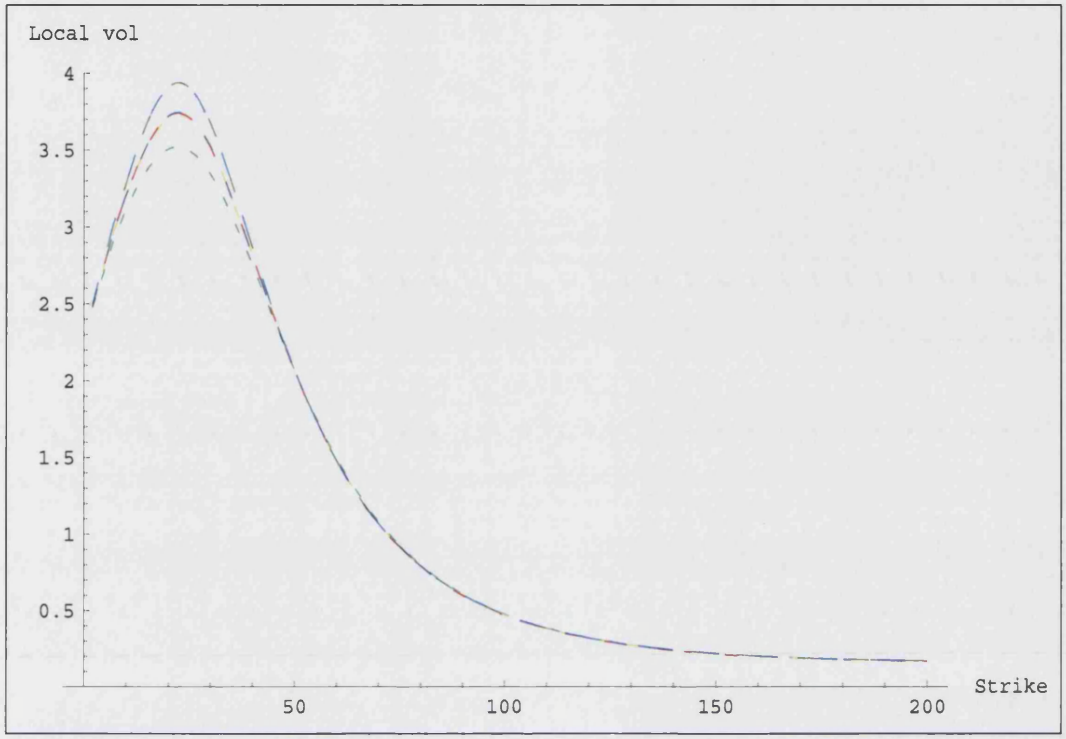

Figure 2.6: Local volatility functions (P1) to (P6) for the skew-case and $T_0 = 1/12$

The corresponding local volatility functions (P1) to (P6) are shown in Figure 2.6 for the skew case and Figure 2.7 for the smile case. In these figures, and also in the remainder of this chapter, we use cyan for (P1), yellow for (P2), lilac for (P3), red for (P4), green for (P5) and blue for (P6).

As we can see, the local volatility functions exhibit pronounced non-monotonicity for low strikes, i.e. in a region where the corresponding implied volatility functions are monotonic. The slopes and levels of the local volatilities are even more extreme than those of the underlying local volatilities; while the implied volatility skew does not exceed 120%, the local volatility almost attains a level of 400%. We observe that the local volatility functions are almost identical for strikes greater than 50, while they show a marked difference for lower strikes. This suggests a high degree of time-inhomogeneity in the exact local volatility function in this region. Furthermore, it is striking that methods (P2) to (P4) and (P5) and (P6) produce almost identical local volatilities throughout the whole strike range. We will observe this phenomenon in our other scenarios as well.

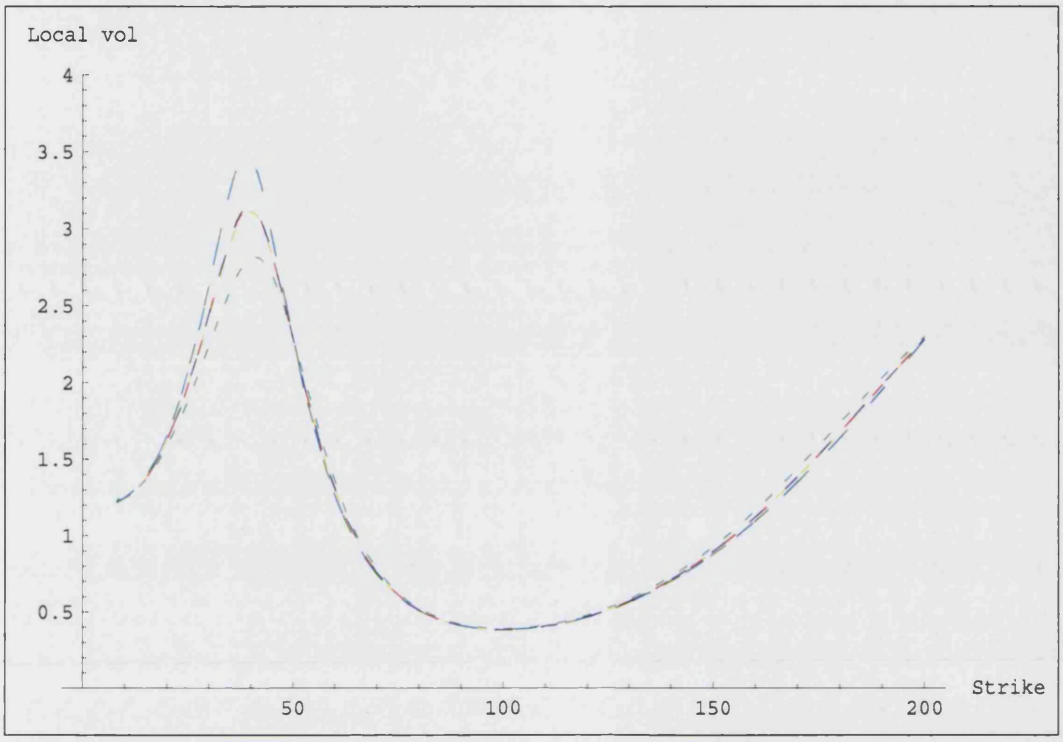

Figure 2.7: Local volatility functions (P1) to (P6) for the smile-case and $T_0 = 1/12$

Comparing the exact implied volatilities (dotted black lines) with the model implied volatilities (cf. Figures 2.8 and 2.9), and the corresponding approximation errors, i.e. the differences between model- and exact implied volatilities (cf. Figures 2.10 and 2.11), we find a remarkably good fit of the model implied volatilities to the given ones. While even for the simple approximation (P1) the approximation errors do not exceed 0.4 percentage points for the skew-case and 1.2 percentage points for the smile-case, the more sophisticated approximations (P2) to (P6) perform even better, as the errors never exceed 0.5 percentage points over the whole strike range from 50 to 150. Near the money, i.e. for strikes ranging from 80 to 120, the errors are typically even less than 0.1 percentage points for the methods (P2) to (P6). The corresponding price differences are reported in Appendix A. Finally, it is also worth noting that, although there is a plainly visible difference in the local volatility functions (especially for lower strikes), this is not the case for the implied volatility functions, the reason being that the local volatility functions differ only in regions away from the money. As the impact of local volatilities away from the money

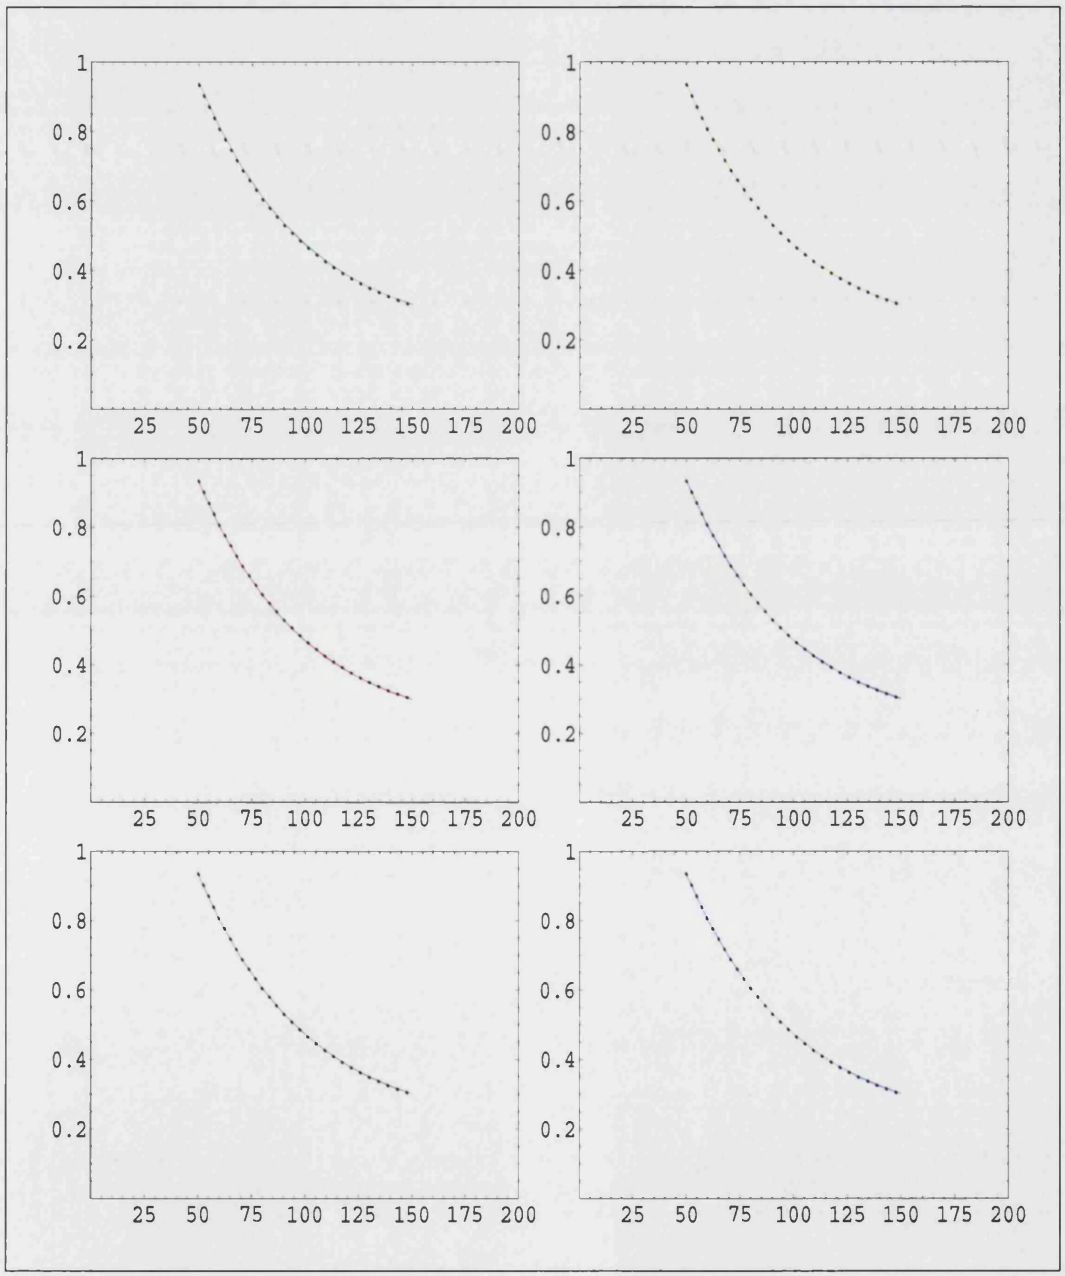

Figure 2.8: Implied volatility functions for (P1) to (P6) for the skew-case and $T_0 = 1/12$

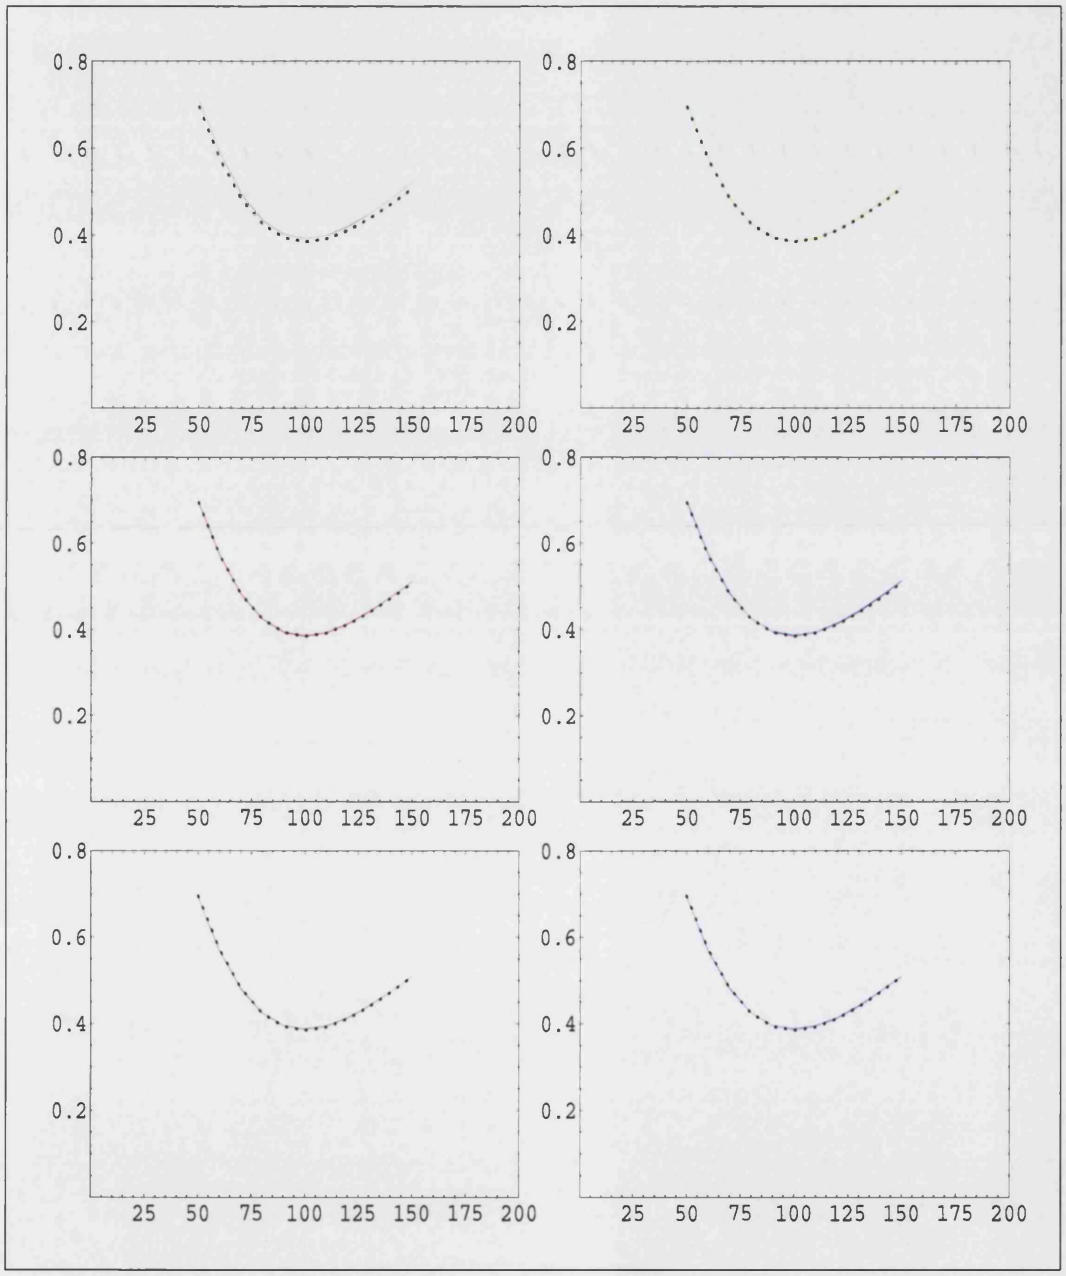

Figure 2.9: Implied volatility functions for (P1) to (P6) for the smile-case and $T_0 = 1/12$

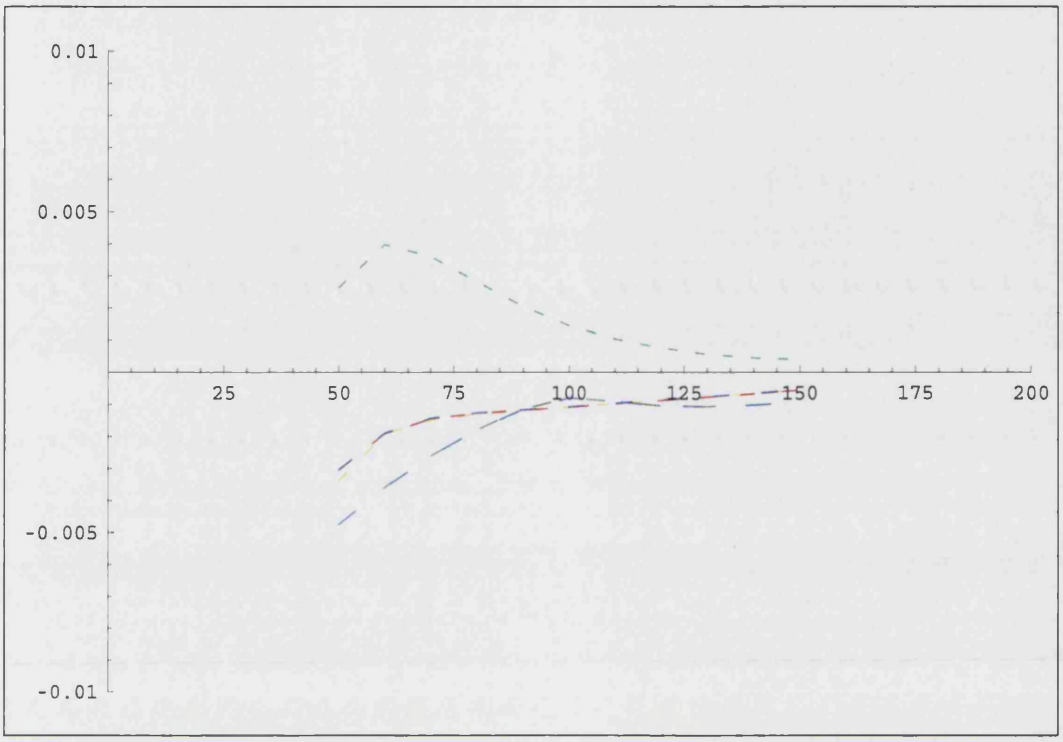

Figure 2.10: Approximation error for (P1) to (P6) for the skew-case and $T_0 = 1/12$

– especially for short maturities – on option prices (respectively implied volatilities) is limited, these differences in local volatilities do not translate into significant differences in implied volatilities.

Medium time to maturity: $T_0 = 1$

Figure 2.12 displays our test functions for $T_0 = 1$, which – though being less extreme in nature – share the characteristics of those for $T_0 = 1/12$. It would not have been possible to use the same implied volatility functions for $T_0 = 1$ as we did for $T_0 = 1/12$, because the corresponding option prices would not have been arbitrage-free for $T_0 = 1$. Although the skew and smile are less pronounced than before, they are also 'extreme' in the sense that the risk-neutral distributions they imply show considerable deviations from

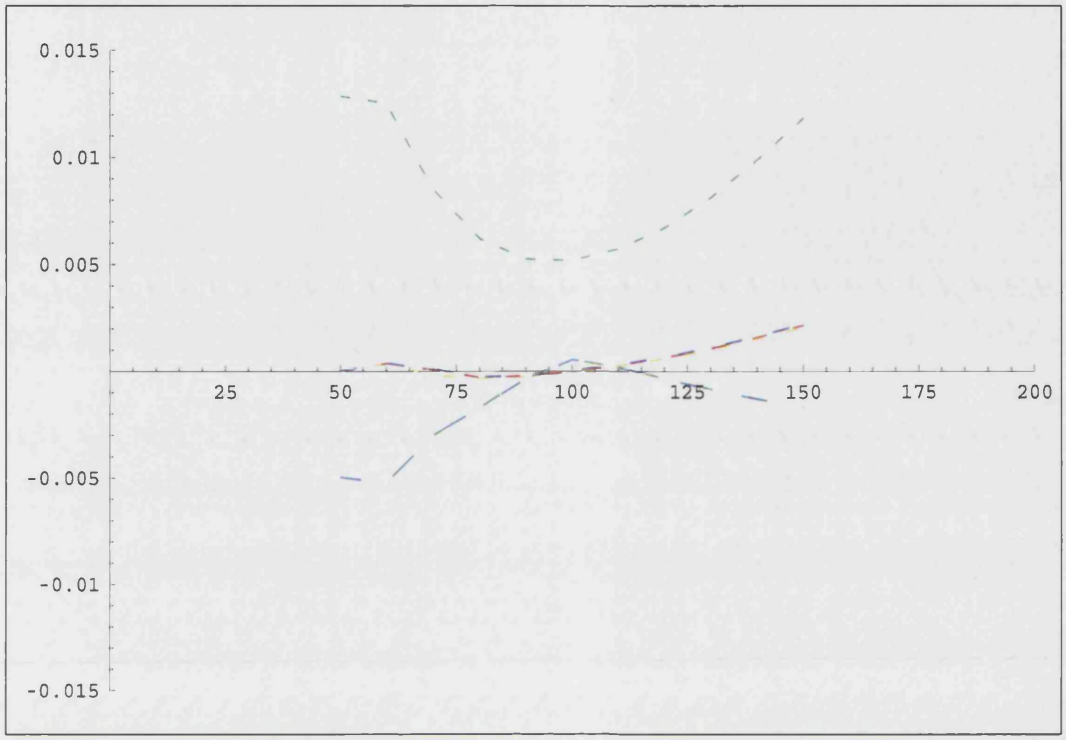

Figure 2.11: Approximation error for (P1) to (P6) for the smile-case and $T_0 = 1/12$

comparable log-normal ones.[17]

Not surprisingly, the corresponding local volatility functions (cf. Figures 2.13 and 2.14), though more moderate, are qualitatively similar to those presented for $T_0 = 1/12$. The main differences compared to the short time-to-maturity case are the smaller first and second derivatives of the implied volatility functions. However, this does not necessarily imply a lesser degree of time-inhomogeneity, because although the functions $u_3(K)$ and $u_4(K)$ (which influence the degree of inhomogeneity) are smaller than in the previous case, they receive higher weights (cf. formula (2.35)), so that we have two counteracting effects, and it is not obvious which one prevails for what K.

As before, we represent the model implied volatilities and the approximation errors graphically (cf. Figures 2.15 to 2.18). To take the greater time to maturity compared to the

[17] In order to get an impression of the deviation of the implied risk-neutral distributions from log-normality, one could for example take log-normal distributions based on the same at-the-money volatilities as a benchmark.

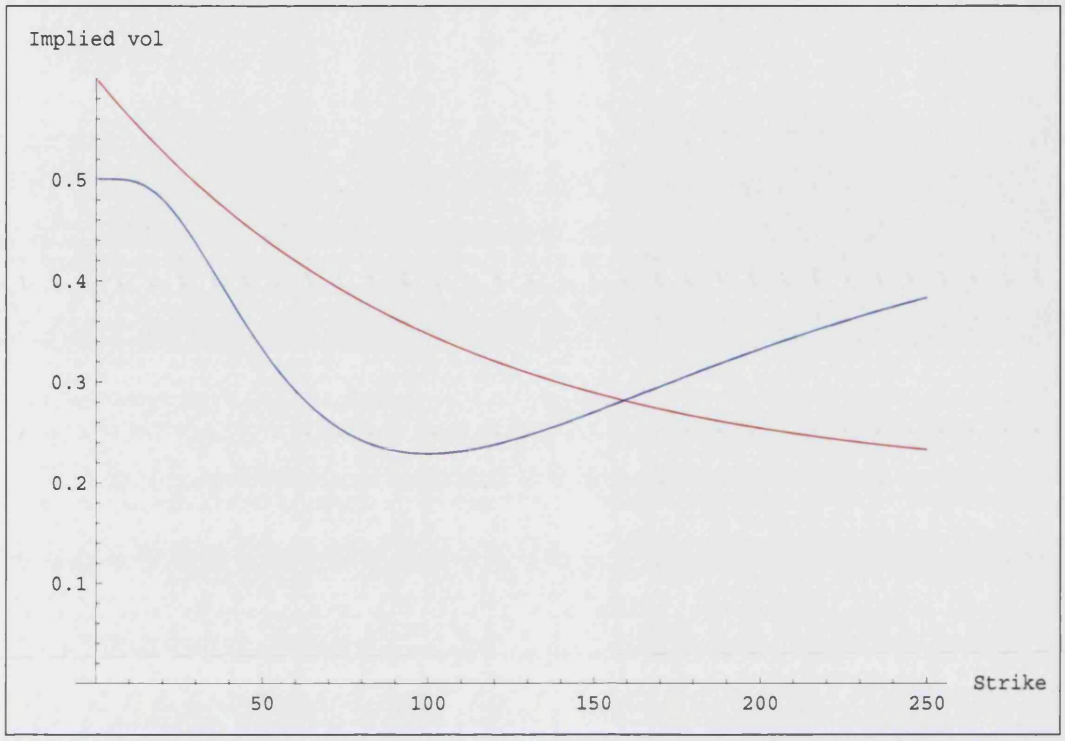

Figure 2.12: Implied volatility skew (red) and smile (blue) for $T_0 = 1$

previous case into account (implying that also implied volatilities further away from the money are significant in terms of information content, and that stable numerical inversion is possible for a broader range of strikes), the curves cover strikes ranging from 40 to 200.[18] Again, we can observe an almost perfect fit to the input volatility structures (dotted black lines) over the whole strike range: The approximation errors for the methods (P2) to (P6) hardly ever exceed 0.1 percentage points for the skew case and 0.2 percentage points for the smile case. The approximation in the skew case seems to be somewhat better than in the smile case. This can be explained by the higher absolute values of the first and second derivatives for the smile, which lead to a comparatively higher degree of time-inhomogeneity, which in turn causes the approximation error to be higher.

[18]The corresponding numerical results can again be found in Appendix A.

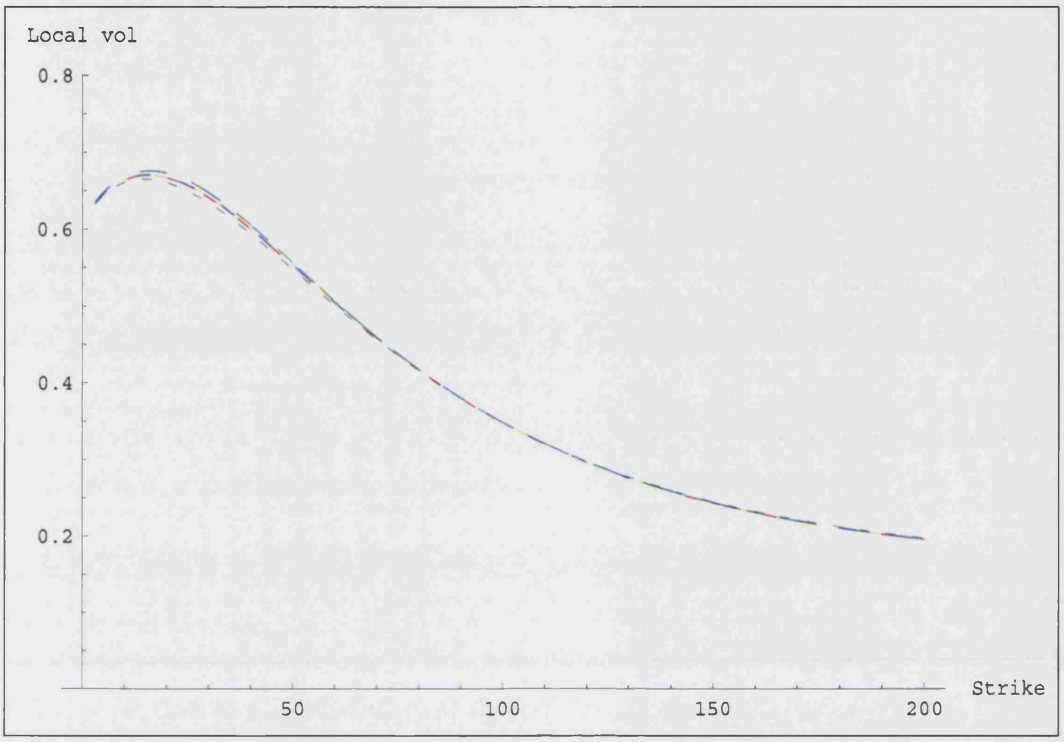

Figure 2.13: Local volatility functions (P1) to (P6) for the skew-case and $T_0 = 1$

Long time to maturity: $T_0 = 10$

Finally, we consider the case $T_0 = 10$. The implied volatility functions that serve as inputs for our tests are depicted in Figure 2.19. Prima facie, they might seem very moderate, but indeed they are not. The long time to maturity imposes strong restrictions on the possible shapes of implied volatility structures, as we demand that the latter be arbitrage-free. But for long maturities, even seemingly moderate implied volatility structures can give rise to risk-neutral distributions that are rather pathological. For example, the risk-neutral distribution implied by our smile is bimodal. One aspect that makes long times to maturity interesting for our purposes is the fact that the terms $u_3(K)$ and $u_4(K)$ in the denominator of the exact local volatility functions are scaled with factors of up to 10 and 100, respectively, suggesting a considerable degree of time-inhomogeneity, even though $u_3(K)$ and $u_4(K)$ are smaller than in the previous cases. Thus we can expect (P1) to perform worse than in the other cases, as it totally neglects these two terms. The local volatility functions in Figures 2.20 and 2.21 look as expected, the only difference being

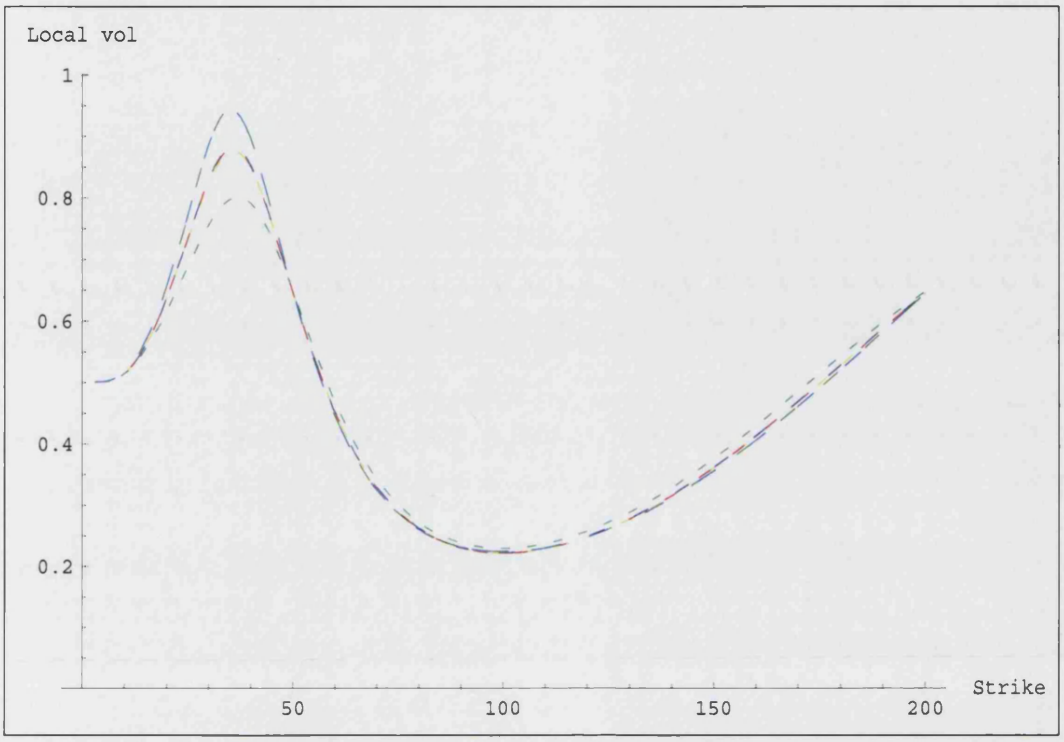

Figure 2.14: Local volatility functions (P1) to (P6) for the smile-case and $T_0 = 1$

that in the smile-case, (P1) now visibly differs from (P2) to (P6) over the whole strike range.

For the implied volatility and error plots (Figures 2.22 to 2.25), we consider a strike range from 30 to 300. Again, as for the shorter maturities, the approximations are very good. The maximum approximation errors for (P2) to (P6) are always less than 0.1 percentage points for the skew and 0.5 percentage points for the smile. Even the rather naïve approximation (P1) never differs by more than 0.3 percentage points for the skew and 1 percentage point for the smile from the exact value.

Summary

The above six test scenarios were chosen so as to cover both skews and smiles, which are the volatility patterns most often encountered in practice, and a maturity spectrum ranging from very short to very long maturities. Summarising our test results, we can claim that our approximations did very well throughout all scenarios, which are characterised by

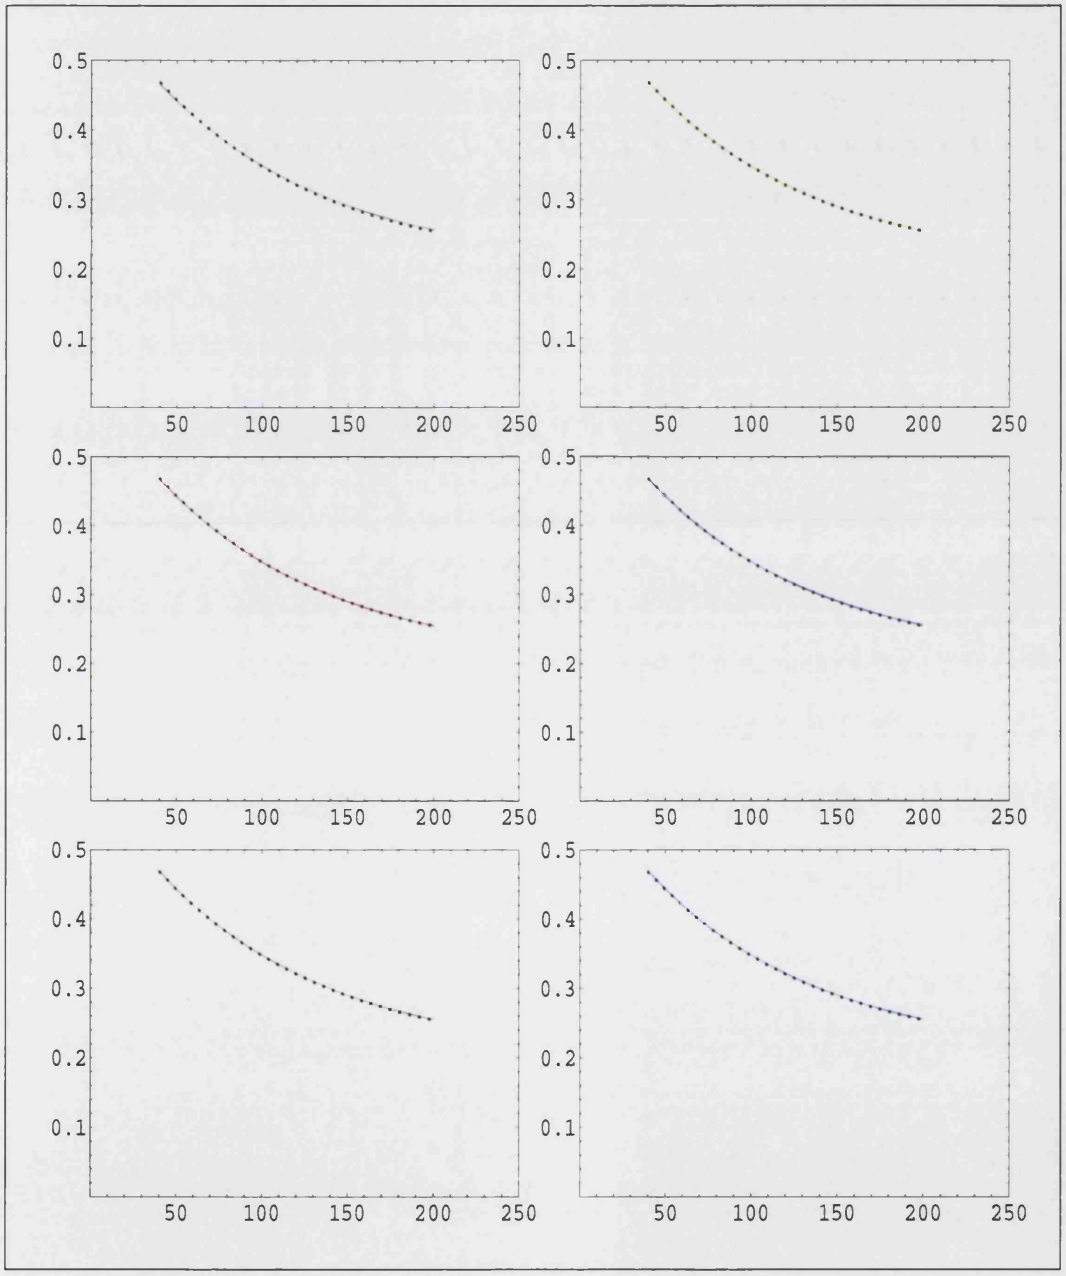

Figure 2.15: Implied volatility functions for (P1) to (P6) for the skew-case and $T_0 = 1$

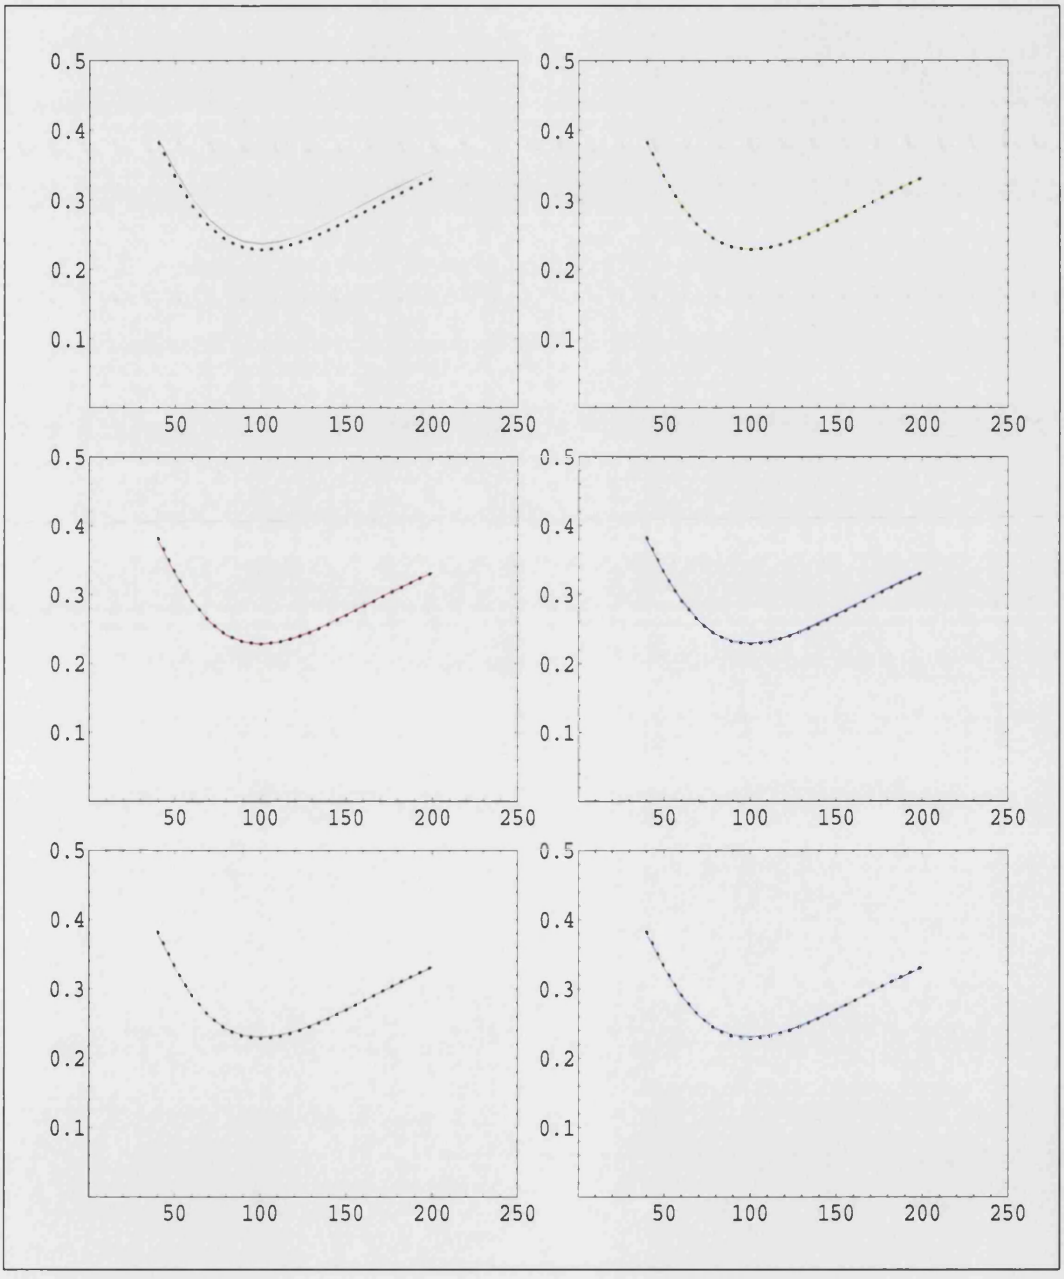

Figure 2.16: Implied volatility functions for (P1) to (P6) for the smile-case and $T_0 = 1$

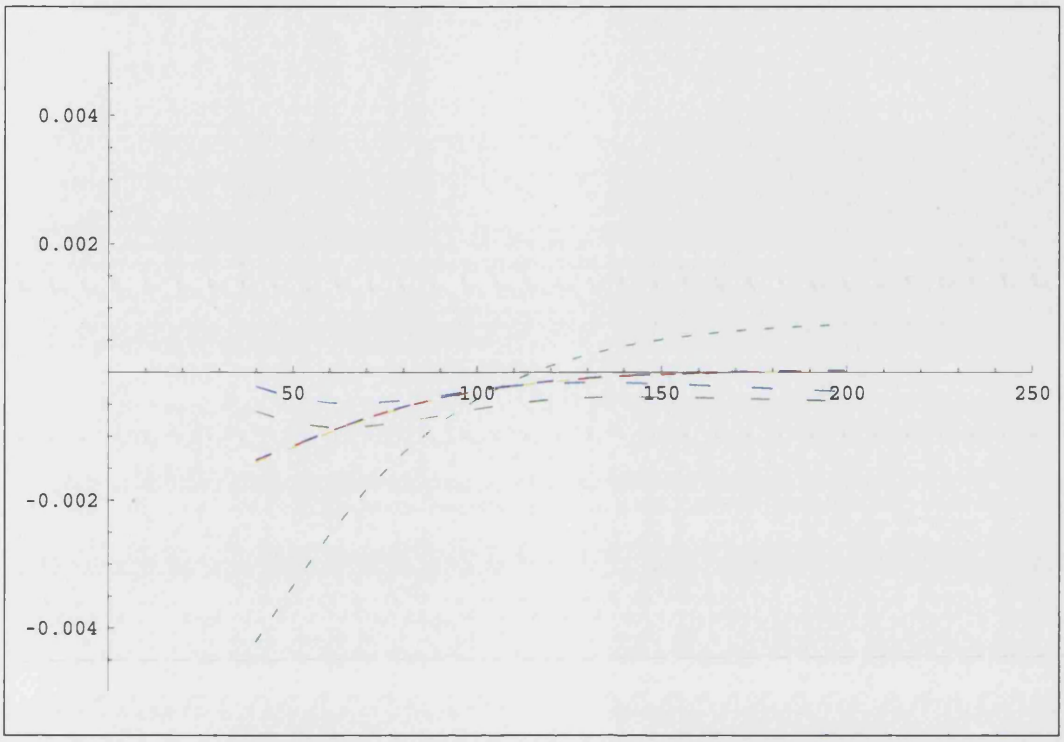

Figure 2.17: Approximation error for (P1) to (P6) for the skew-case and $T_0 = 1$

rather extreme volatility patterns. Therefore, we can hope that they perform even better in most real-world cases, which are likely to be more moderate. A comparison of the quality of the individual approximations confirms our initial conjecture that methods (P2) to (P6) outperform (P1). However, it is not possible to identify a single best approximation that is superior to all others throughout. Perhaps surprisingly, the rather simple method (P2) seems to be of similar quality than the more complex methods (P3) to (P6). As we already remarked above, (P5) and (P6) have to be worked out by means of numerical integration and are therefore computationally quite expensive and prone to numerical errors. This leads us to conclude that, while (P1) is by far sufficient for most practical applications, (P2) to (P4) are the methods of choice when accuracy is an issue, as they strike a good balance between computational complexity and goodness of approximation.

Before concluding this section, it is worth mentioning that the quality of the approximations is invariant under scaling, i.e. independent of the absolute numerical values of the prices used, the reason being that we could have also worked in terms of relative moneyness

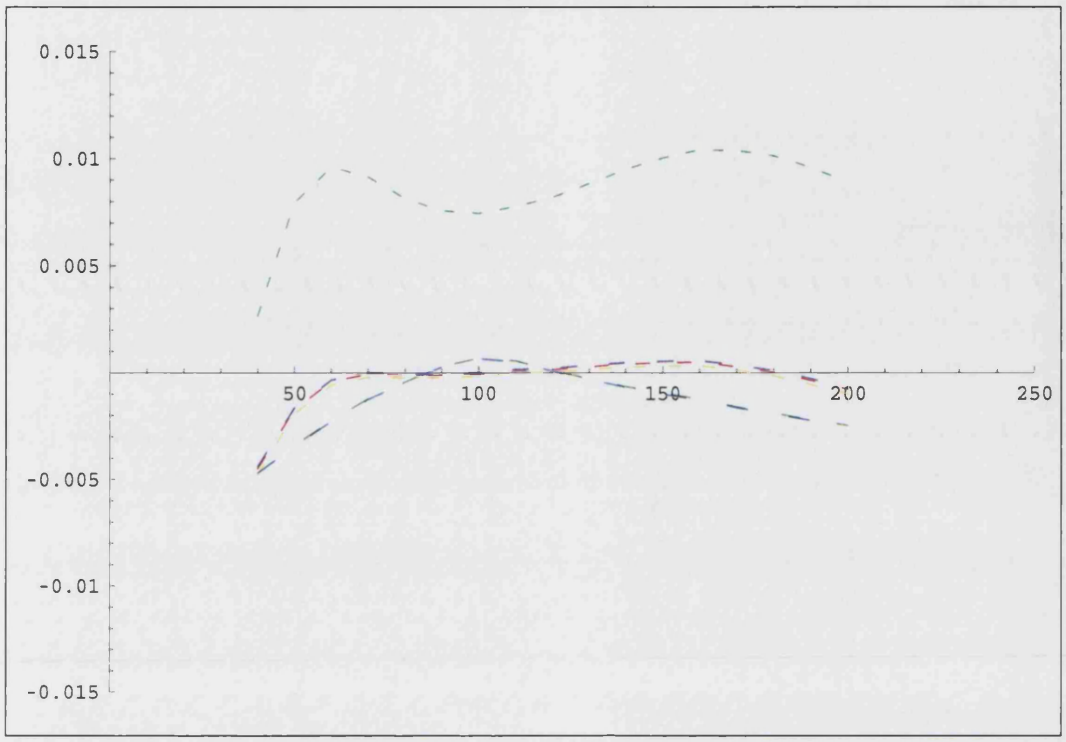

Figure 2.18: Approximation error for (P1) to (P6) for the smile-case and $T_0 = 1$

F/K instead of using absolute values.

2.5 Local volatility models in practice

At this point, the reader might ask why we bother with time-homogeneous (i.e. purely level-dependent) local volatility functions that are in most cases only approximate solutions to the single smile problem, even though we are able to state time-dependent local volatility functions that are exact solutions (compare formula (2.35)) or even fit a whole given volatility surface. The main reasons are of practical nature: In practice, one only encounters a finite, discrete set of vanilla prices that can serve to calibrate the model. It might well be the case that there is only one option-maturity for which the corresponding options are liquid enough to reliably serve as calibration instruments, so that the calibration-problem in practice often boils down to the single smile problem. Even if there exist a couple of liquid maturities, it might not be desirable to incorporate all of them

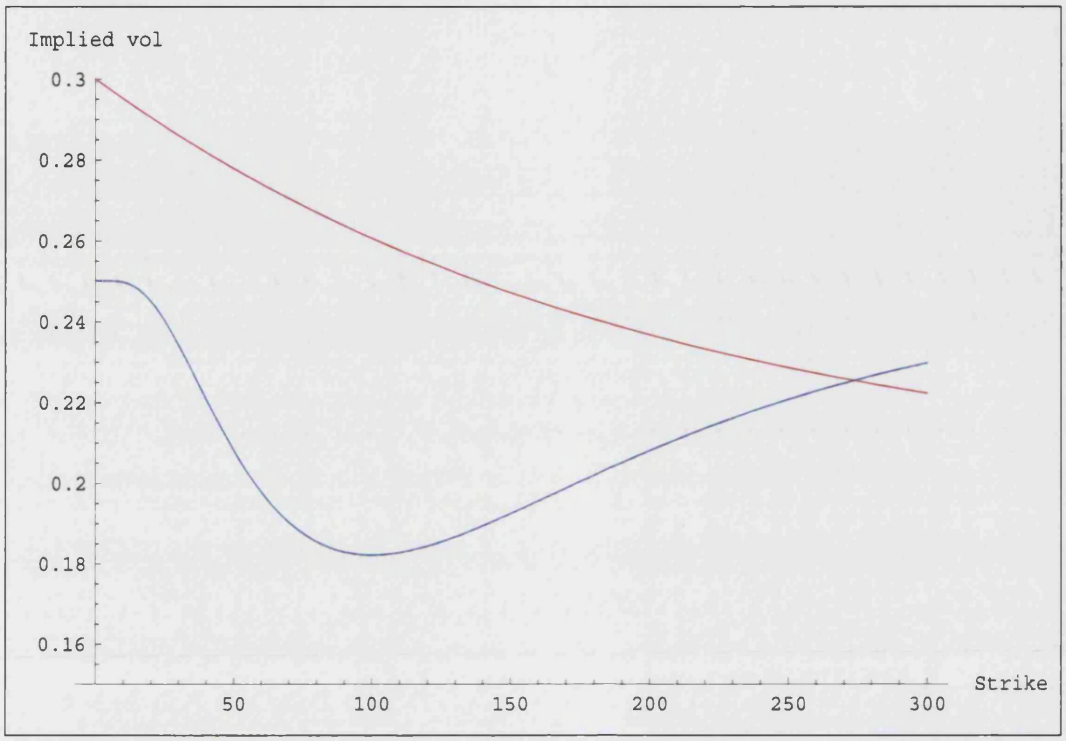

Figure 2.19: Implied volatility skew (red) and smile (blue) for $T_0 = 10$

in the calibration process, as this often leads to wildly fluctuating and implausible local volatility functions. However, a time-homogeneous local volatility function, appropriately scaled with a purely time-dependent function, leads to a separable volatility function that gives the user full control over the term structure of local and implied volatilities. This is not the case with non-separable local volatility functions obtained by the classical Dupire approach. In addition, there exists a very accurate analytical approximation for option prices if the dynamics of the underlying are governed by a separable local volatility function, whereas for non-separable ones, one generally has to resort to numerical methods. We will now elaborate on the above points.

2.5.1 Separable local volatility functions

A separable local volatility function can be factorised: It can be represented as $\lambda(T)\sigma(K)$ with a bounded and deterministic function $\lambda : \mathbb{R}^+ \mapsto \mathbb{R}^+$. Now assume that we are given a purely level-dependent volatility function $\sigma(K)$, which for example could be obtained

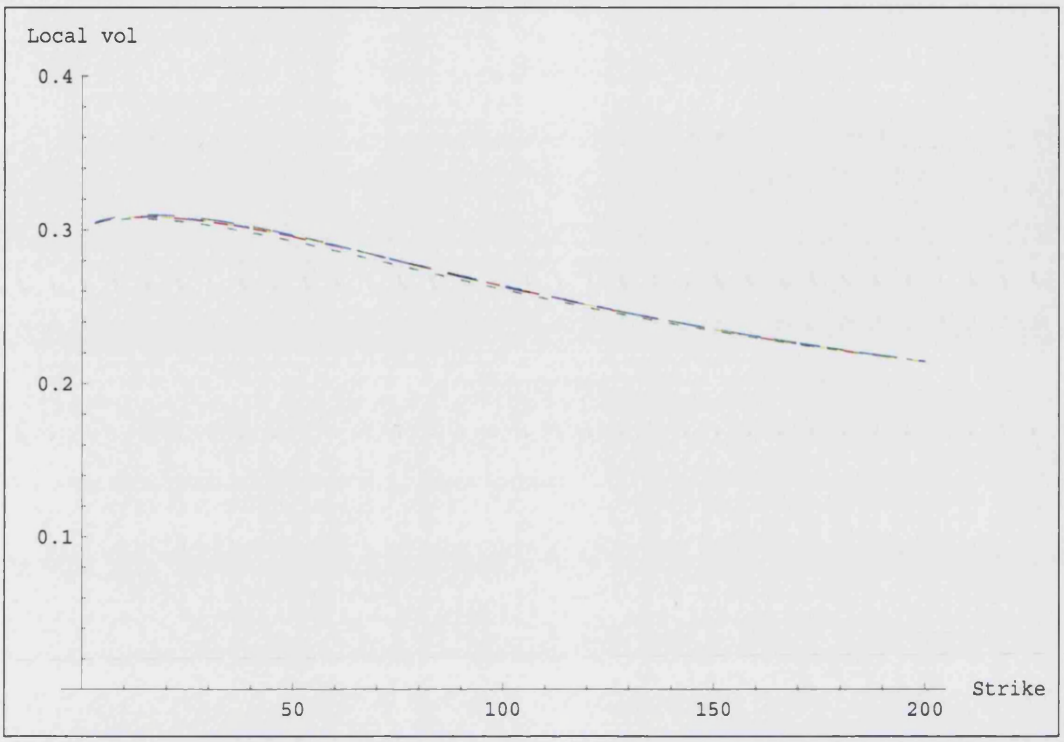

Figure 2.20: Local volatility functions (P1) to (P6) for the skew-case and $T_0 = 10$

by the methodologies detailed in the previous section. We describe the term structure of local volatility between $t = 0$ and T_0 by a function $\lambda : [0, T_0] \mapsto \mathbb{R}^+$ with the property

$$\int_0^{T_0} \lambda(T)^2 \, dT = T_0.$$

Suppose that the T_0-forward price follows

$$dF(T, T_0) = \lambda(T)\sigma(F(T, T_0)) \, dW_T. \tag{2.36}$$

To keep our notation simple, we set

$$v(T) = \int_0^T \lambda(s)^2 \, ds.$$

Then, by a standard deterministic time-change argument (cf. e.g. Øksendal [2000], Chapter 8.5), $\left(\tilde{W}_T\right)$ with

$$\tilde{W}_{v(T)} = \int_0^T \lambda(s) \, dW_s$$

is a standard Brownian motion. Setting $f(v(T)) = F(T, T_0)$, we get

$$df(v(T)) = dF(T, T_0) = \lambda(T)\sigma(F(T, T_0)) \, dW_T = \sigma(f(v(T)) \, d\tilde{W}_{v(T)},$$

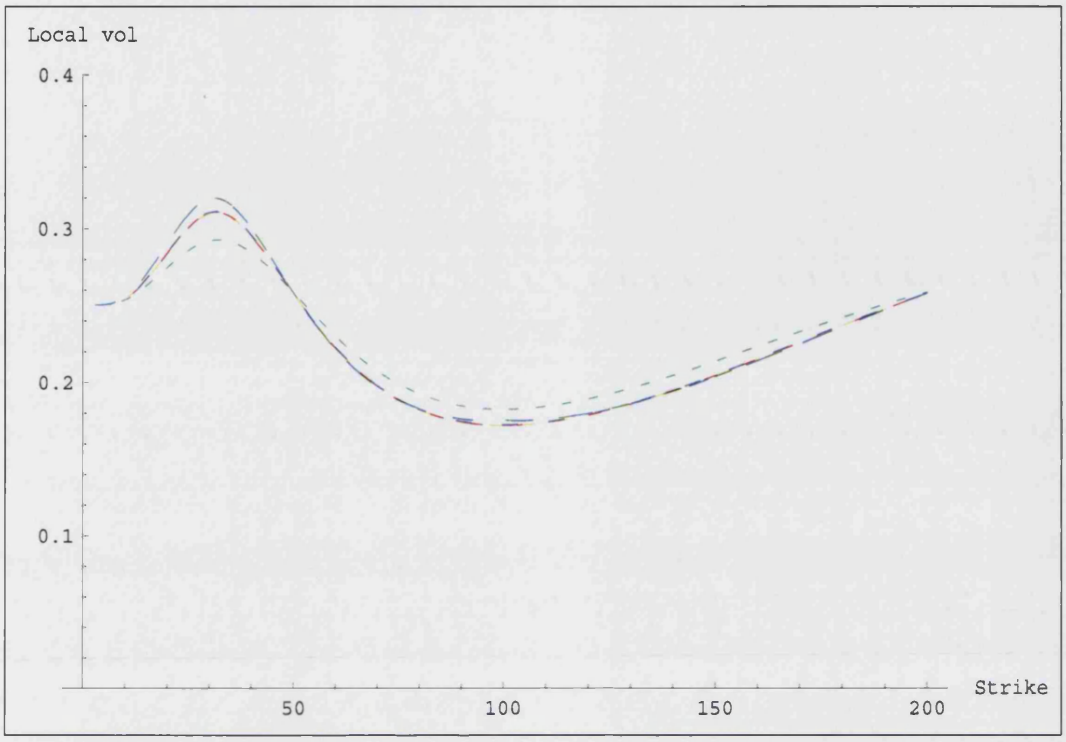

Figure 2.21: Local volatility functions (P1) to (P6) for the smile-case and $T_0 = 10$

or equivalently, as integral equation,

$$f(v) = \int_0^v \sigma(f(u))\, d\tilde{W}_u.$$

Now,

$$F(T_0, T_0) = f(v(T_0)) = f(T_0) = \int_0^{T_0} \sigma(f(u))\, d\tilde{W}_u,$$

but also

$$F(T_0, T_0) = \int_0^{T_0} \lambda(T)\sigma(F(T, T_0))dW_T.$$

This demonstrates that multiplying by an appropriately scaled function λ does not alter the distributional properties of $F(T_0, T_0)$, while allowing the user to exogenously specify a volatility term structure. As many exotic options (e.g. forward-start options) are very sensitive to the term structure of local volatility, the specification of a realistic λ is of prime importance, as it will severely affect the pricing and hedging performance of the model.

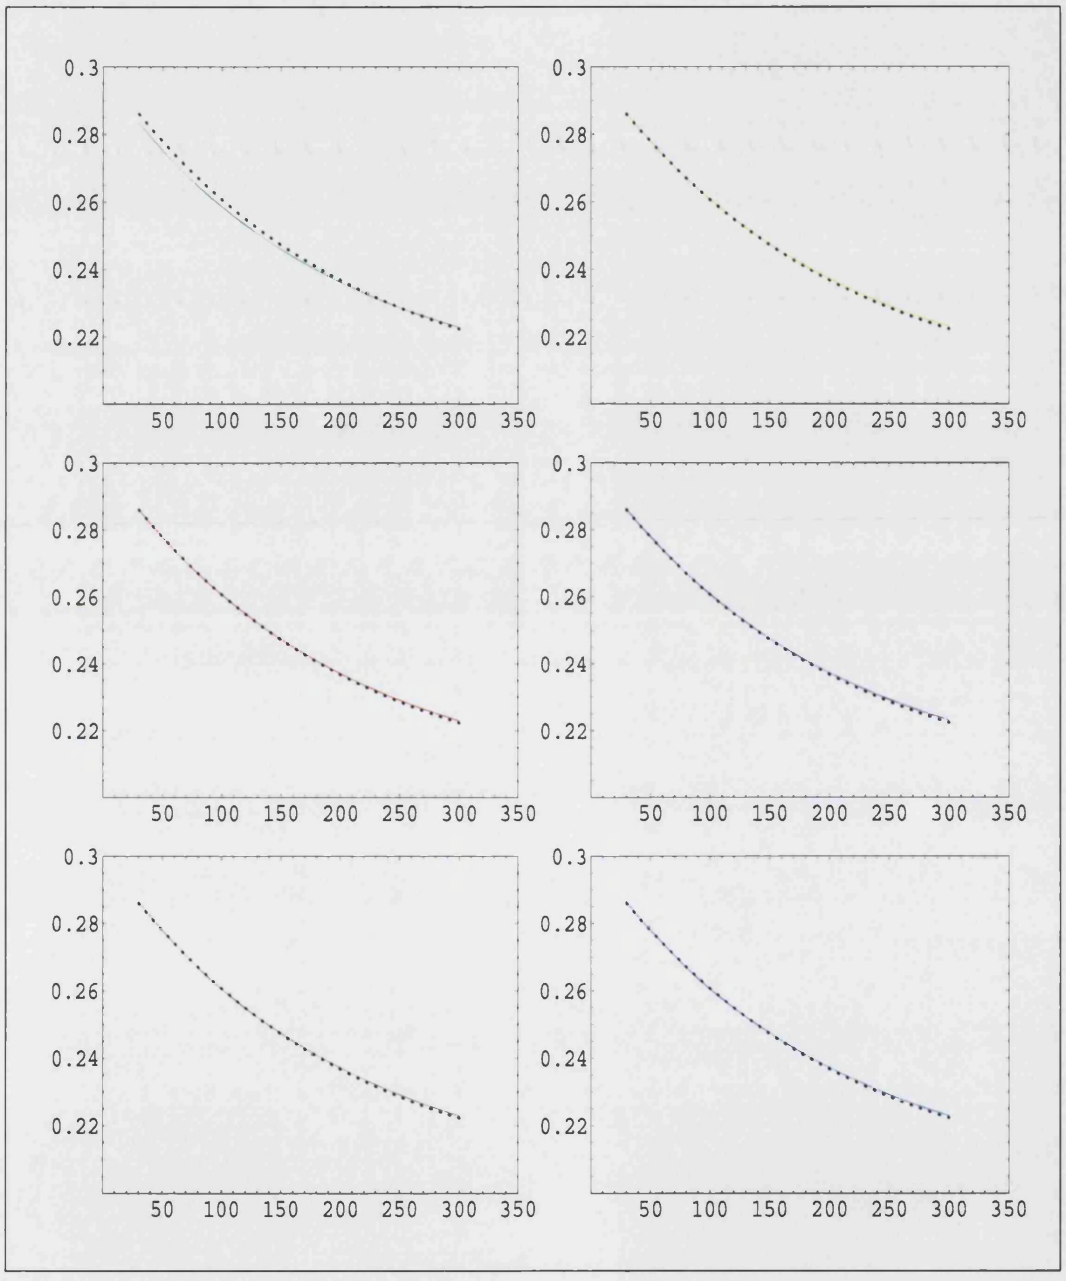

Figure 2.22: Implied volatility functions for (P1) to (P6) for the skew-case and $T_0 = 10$

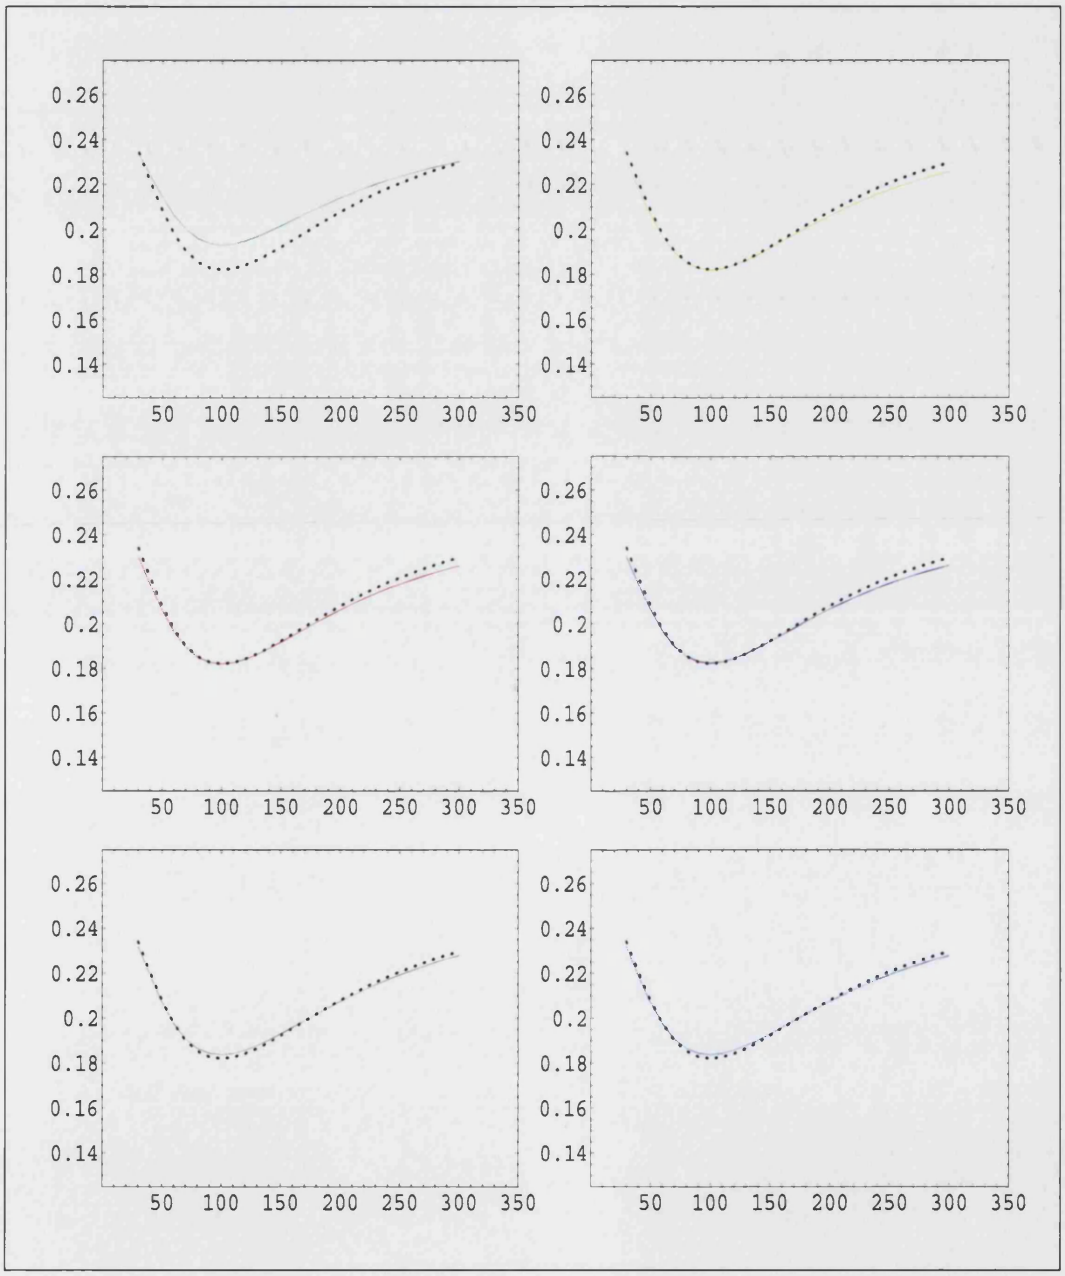

Figure 2.23: Implied volatility functions for (P1) to (P6) for the smile-case and $T_0 = 10$

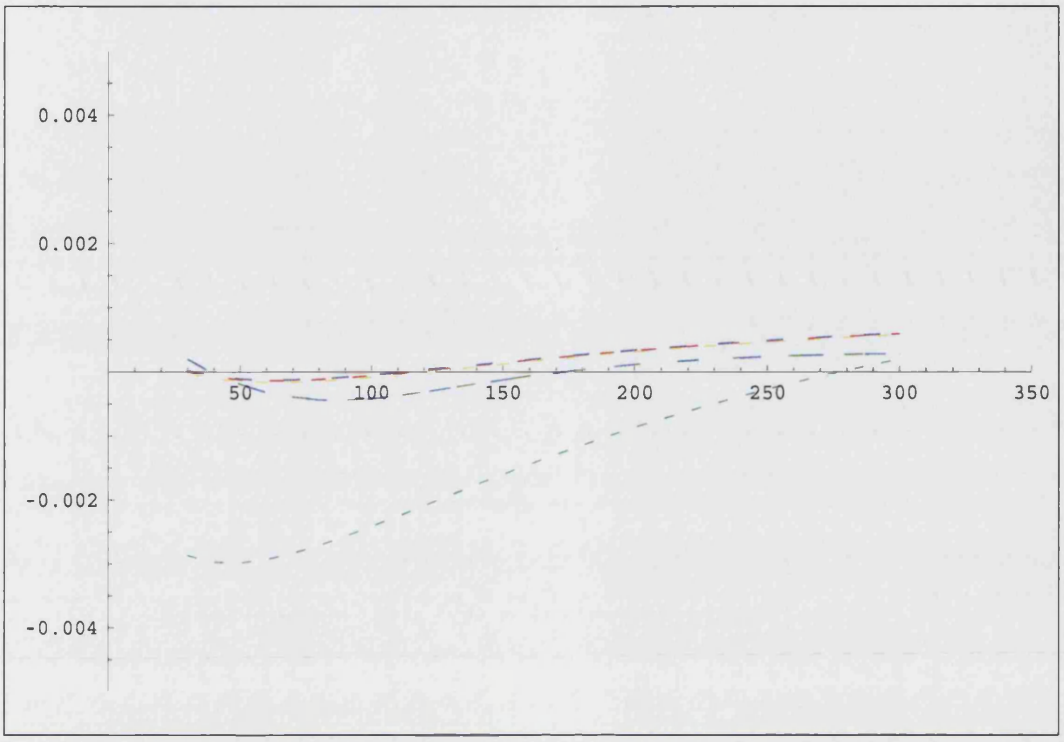

Figure 2.24: Approximation error for (P1) to (P6) for the skew-case and $T_0 = 10$

2.5.2 Analytical approximations

Apart from the ability to fully control the term structure of volatility while matching a future smile, separable local volatility functions offer another important advantage: While for non-separable local volatility functions, even standard options have to be priced by numerical methods, there exist very good analytical approximations for European calls and puts for separable ones. Hagan & Woodward [1999] propose the following approximation. Suppose that $(F(T,T_0))$ follows (2.36). Then the $t=0$-price $C^F(0,T,K)$ of a European call option with maturity $T < T_0$ and strike K on the T_0-forward price forward price can be approximated by $C^{Black}(F(0,T_0),0,T,K,\sigma^{imp},(r_s))$ with

$$\sigma^{imp} = \frac{aA(F_{av})}{F_{av}}\left[1+\frac{1}{24}\left(\frac{A''(F_{av})}{A(F_{av})}-2\frac{A'(F_{av})^2}{A(F_{av})^2}+\frac{2}{F_{av}^2}\right)(F(0,T_0)-K)^2\right.$$
$$\left.+\frac{1}{24}\left(2\frac{A''(F_{av})}{A(F_{av})}-\frac{A'(F_{av})^2}{A(F_{av})^2}+\frac{1}{F_{av}^2}\right)a^2A(F_{av})^2T\right], \tag{2.37}$$

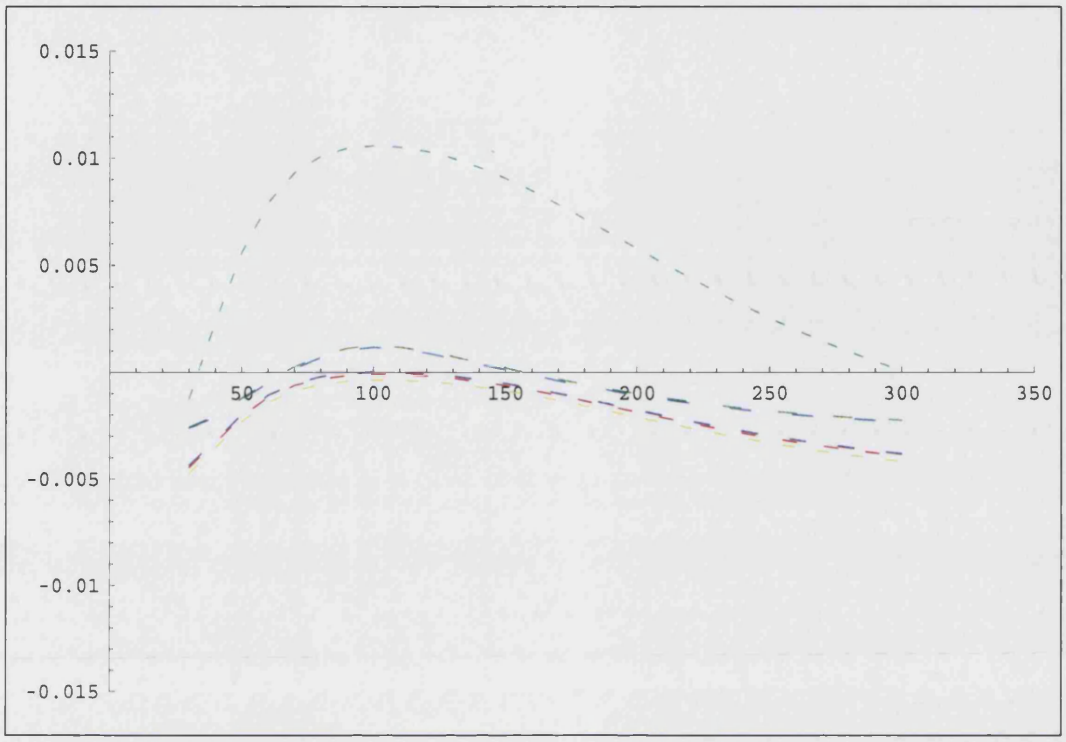

Figure 2.25: Approximation error for (P1) to (P6) for the smile-case and $T_0 = 10$

where

$$a = \left(\frac{1}{T} \int_0^T \lambda(s)^2 \, ds \right)^{\frac{1}{2}}$$

and

$$F_{av} = \frac{F(0,T_0) + K}{2}.$$

Hagan and Woodward stress the excellent accuracy of their approximation, which they show to be of similar quality to PDE-based methods.

2.5.3 Hedging in local volatility models

Local volatility models are complete one-factor models. This implies that any derivative can be perfectly replicated. For hedging- or replication purposes, it is important to be able to calculate the Greeks of an option – in particular its delta – quickly and accurately.

The delta of a European call is

$$\begin{aligned}
&\frac{\partial}{\partial F} C^{Black}(F(t,T_0),t,T,K,{}_t\sigma^{imp,F}(F(t,T_0),T,K),(r_s)) \\
&= \Delta^{Black}(F(t,T_0),t,T,K,{}_t\sigma^{imp,F}(F(t,T_0),T,K),(r_s)) \\
&\quad + \mathcal{V}^{Black}(F(t,T_0),t,T,K,{}_t\sigma^{imp,F}(F(t,T_0),T,K),(r_s)) \cdot {}_t\sigma_F^{imp,F}(F(t,T_0),T,K),
\end{aligned} \tag{2.38}$$

where Δ^{Black} and $\mathcal{V}^{Black}$ denote the standard delta and vega in the Black model. Here, it should be emphasised that ${}_t\sigma^{imp,F}$ is determined by the a priori given local volatility function and therefore endogenous to the model. As already mentioned, in a local volatility setting, option prices (and also implied volatilities) normally have to be worked out by numerical methods, and this apparently also applies to the calculation of delta. But approximation (2.37) can also be utilised to calculate approximations for delta, which can also be expected to be highly accurate.

Formula (2.38) indicates that the implied volatility of an option is liable to change as time unfolds or as the price of the underlying changes.[19] This raises the issue of implied volatility dynamics in a local volatility model. The specification of a local volatility function already fully determines the dynamics of the whole implied volatility surface (i.e. the implied volatility surface that extends from the current time to the end of the modelling horizon and over all positive K): A forward price $F(T,T_0)$ observed at a future time T is uniquely linked to the implied volatility surface that prevails in the state $(T,F(T,T_0))$.[20] A natural question is whether future implied volatility surfaces generated by local volatility models are realistic. There is a consensus in the literature that, unfortunately, local volatility models predict the wrong dynamics of the implied volatility surface (see e.g. Hagan *et al.* [2002]), as they result in implied volatility skews/smiles moving in the opposite direction to the price of the underlying, which is in contrast to typical market behaviour, where skews/smiles and underlying move in the same direction. This is believed to be the main reason for inaccurate and unstable hedges often encountered in a local volatility

[19] A delta hedging strategy according to (2.38) also implicitly provides a hedge against changes in implied volatility due to changes in the price of the underlying.

[20] This is one of the major differences between local- and stochastic volatility models: In a stochastic volatility model, different implied volatility surfaces can prevail in a future state $(T,F(T,T_0))$.

framework, as reported by Dumas *et al.* [1998] and Balland [2002]. It is also this deficiency that has spurred much of the recent interest in stochastic implied volatility models (see e.g. Brace *et al.* [2001], Brace *et al.* [2002], Ledoit *et al.* [2002] and Daglish *et al.* [2003]), which can be perfectly calibrated to observed implied volatility patterns, while still granting relative freedom when it comes to exogenously specifying implied volatility dynamics.

2.6 Summary

In this chapter, we present a methodology for constructing local volatility functions that exactly reproduce given volatility smiles. In this context, we give explicit analytical formulae for state-price densities and transition densities. Furthermore, we propose approximate solutions to the single smile problem: Given an implied volatility function, we are able to state a time-homogeneous local volatility function that approximately reproduces the given implied volatility function. Numerical tests based on several extreme volatility scenarios show that our methods provide an excellent fit to the input implied volatility structure over a wide range of strikes and maturities. Being based directly on implied volatilities rather than option prices, our method avoids the problem of interpolating option prices, as it is by no means clear how to interpolate a discrete set of option prices in such a way that the radicand in the classical Dupire formula is well-defined (see Berestycki *et al.* [2002]). In contrast, direct interpolation of implied volatilities is more robust.

The methods detailed in this chapter are tailor-made for the construction of smile-consistent LIBOR market models, which is the content of the next chapter.

Chapter 3

Smile-consistent generalised extended LIBOR market models

The so-called *market models* of LIBOR and swap rates have enjoyed increasing popularity during the last few years. The main reason for the success of this approach to term structure modelling, developed in a series of papers by Brace *et al.* [1996], Miltersen *et al.* [1997] and Jamshidian [1997], can be seen in its consistency with the market practice of pricing caps, floors and swaptions by means of the Black [1976]-formula, while at the same time providing a consistent and coherent framework for the joint modelling of a whole set of forward rates. Further aspects that set the class of market models apart from other interest rate models is the relative ease of calibration to market data (e.g. term structures or cap-prices), and the use of *discretely* compounded forward LIBOR rates – directly observable in the market – as fundamental quantities in the modelling process. By contrast, traditional short-rate or Heath-Jarrow-Morton models are based on the description of the arbitrage-free dynamics of *continuously* compounded instantaneous short or forward rates, which are not market observables.

Recently, several studies (e.g. Andersen & Andreasen [2000], Joshi & Rebonato [2001], Joshi & Rebonato [2003] and Jarrow *et al.* [2003]) have documented the presence of non-flat IVS (when quoted in terms of Black implied volatilities) in cap and floor markets. The inability of standard LIBOR market models (LMMs for short) to capture these smiles

or skews has spurred the development of the class of extended LIBOR market models (henceforth EMMs), which we will discuss in the following, where we will find that even models of this class – although more general than standard LMMs – are not flexible enough to meet all practical requirements. For this reason, we will extend the existing theory on EMMs to a more flexible model class, which we will term generalised extended LIBOR market models (GEMMs).

This chapter is structured as follows. After introducing the mathematical set-up, we will formally develop the theory of GEMMs. Then, we shall demonstrate how to price caplets and floorlets in the GEMM setting. Subsequently, we will focus on pricing swaptions and derive a swaption-approximation, which will be tested in several scenarios. Finally, we shall touch on some issues concerning the practical implementation of our model, and conclude with a short summary.

3.1 Description of the economy

We start by defining the *tenor structure* $\mathcal{T} = \{T_0, \ldots, T_n\}$ as a set of maturities T_i with $0 = T_0 < T_1 < \cdots < T_n$, where T_n is the time horizon of our economy. A given tenor structure $\mathcal{T}$ is associated with a set of $\{\tau_1, \ldots, \tau_n\}$ of year fractions, where $\tau_i = T_i - T_{i-1}$, $i = 1, \ldots, n$. We assume that in the financial market under consideration, there exist zero-coupon bonds $p(\cdot, T_i)$ of all maturities T_i, $i = 1, \ldots, n$. The discretely compounded *forward LIBOR rate* prevailing at time t over the future period from T_{i-1} to T_i is defined by

$$L(t, T_{i-1}) = \frac{p(t, T_{i-1}) - p(t, T_i)}{\tau_i p(t, T_i)}, \quad 0 \leq t \leq T_{i-1}.$$

We assume further that we are given a stochastic basis $\left(\Omega, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T_n]}, \mathcal{F}, \mathbb{P}\right)$, on which a d-dimensional Brownian motion $W = (W(t))_{t \in [0, T_n]}$ is defined. The stochastic basis is assumed to satisfy the usual conditions.

3.2 From the standard LMMs to GEMMs

We are now in the position to introduce the important notion of the *forward LIBOR measure.*

Definition 3.1. *Let $T_k \in \mathcal{T}$. We call a $\mathbb{P}$-equivalent probability measure $\mathbb{Q}^{T_k}$ the* forward LIBOR measure for the maturity T_k, *or more briefly T_k* forward measure, *if all bond price processes*

$$\left(\frac{p(t,T_i)}{p(t,T_k)}\right)_{t\in[0,\min\{T_i,T_k\}]}, \quad i=1,\dots,n,$$

relative to the numéraire $p(\cdot,T_k)$ follow local martingales under $\mathbb{Q}^{T_k}$. Correspondingly, the tuple $\left(\mathbb{Q}^{T_k}, p(\cdot,T_k)\right)$ is called the numéraire pair for the maturity T_k.

Now – bearing in mind the definition of $L(\cdot,\cdot)$ – it can be easily observed that the process $(L(t,T_{k-1}))_{t\in[0,T_{k-1}]}$ of the forward LIBOR is a local martingale under $\mathbb{Q}^{T_k}$. So if we place ourselves in a diffusion setting, we can posit, for every $k \in \{1,\dots,n\}$, the following driftless dynamics under the respective forward LIBOR measure $\mathbb{Q}^{T_k}$:

$$\begin{aligned} dL(t,T_{k-1}) &= \sigma(t,T_{k-1})dW^{T_k}(t) \\ &= \sum_{i=1}^{d} \sigma_i(t,T_{k-1})\, dW_i^{T_k}(t), \quad 0 \le t \le T_{k-1}, \end{aligned}$$

with a standard d-dimensional Brownian motion W^{T_k} with respect to $\mathbb{F}$ under the measure $\mathbb{Q}^{T_k}$, and a d-dimensional row vector σ as adapted volatility function that satisfies the standard integrability conditions. This general framework allows for a wide variety of interest rate dynamics. In order to derive further results, one has to impose a certain structure on σ. For instance, the standard LIBOR market models mentioned in the introduction are specified by volatility functions of the form

$$\sigma(t,T_k) = L(t,T_k)\lambda(t,T_k) \tag{3.1}$$

with bounded and deterministic functions $\lambda(\cdot,T_i) : [0,T_i] \to \mathbb{R}^d_+$, $T_i \in \{T_1,\dots,T_{n-1}\}$. This formulation leads to log-normally distributed LIBOR-rates, and is therefore consistent with the market practice of pricing caplets with the standard Black-formula. While λ can be used to capture the term structure of volatility (for example to calibrate the

model to the prices of at-the-money caplets of different maturities), the above formulation does not allow for calibration to implied volatility structures observed in the market. As all LIBORs in this formulation are log-normal under their respective forward measure, the model can only reproduce flat (level-independent) volatility structures. This limitation motivates the generalisation to *separable* volatility functions of the form

$$\sigma(t,T_k) = \phi(L(t,T_k))\lambda(t,T_k),$$

$\phi : \mathbb{R}_+ \to \mathbb{R}_+$, which characterise the class of *extended LIBOR market models* (for short EMMs). Within this framework, specialising ϕ further to

$$\phi(x) = x^\alpha, \quad \alpha > 0,$$

or

$$\phi(x) = x \min\left\{x^{\alpha-1}, \epsilon^{\alpha-1}\right\}, \quad \alpha, \epsilon > 0,$$

leads to the classes of *constant elasticity of variance (CEV)* and *limited constant elasticity of variance (LCEV)* models studied in Andersen & Andreasen [2000]. Setting

$$\phi(x) = a + bx + x^2$$

defines the class of *quadratic volatility models* explored in Zühlsdorff [2002]. All of the above formulations can only generate a limited number of different shapes of the implied volatility structure (IVS), as one could have already guessed by the maximum of two parameters available for calibration purposes. The CEV-model, for example, produces downward sloping IVSs for $0 < \alpha < 1$, a flat IVS for $\alpha = 1$, and upward sloping IVSs for $\alpha > 1$. In particular, it is impossible to generate smile-shaped IVSs often observed in the market. A further drawback of EMMs in general is that the function ϕ is independent of T_k, which makes it impossible to fit the model simultaneously to IVSs for multiple maturities T_k. The logical consequence is the introduction of more general functional classes for σ, which are rich enough to allow for a greater variety of IVS shapes. At the far end of this spectrum, we find model-specifications with volatility functions of the general form

$$\sigma(t,T_k) = \hat{\sigma}(t, L(t,T_k), T_k). \tag{3.2}$$

Most notably, the mixture-of-log-normals model introduced in Brigo *et al.* [2003] falls into this class. While this model features closed-form transition densities and caplet-price formulae, and offers the user an arbitrary number of parameters (which depends on the number of log-normal distributions superimposed) for calibration purposes, is it only capable of generating a limited number of different IVS patterns. However, this limitation is specific to the mixture model. In general, formulation (3.2) can accommodate *any* arbitrage-free IVSs. This follows directly from the results of the previous chapter: Using formula (2.35), we are able to specify a volatility function σ that is compatible with any finite set of arbitrage-free IVSs.[1] However, when using formula (2.35), and also in the formulations offered by Brigo & Mercurio [2003], the volatility function σ is completely determined by the IVSs one wants to fit, leaving the user no influence whatsoever on its term-structure. As a consequence, future IVSs induced by the model may be unrealistic and implausible. As for example Balland [2002], Brigo *et al.* [2003] and Rebonato [2004] point out, future IVSs are an important criterion in judging the quality of a model, since they may have a strong impact on its pricing- and hedging-performance. From this point of view, formulation (3.2) is unsatisfactory. For this reason, building on results of the previous chapter, we consider a formulation that is general enough to yield a very good fit to any finite set of IVSs, yet allows the user to retain full control over the volatility term-structure and therefore the evolution of future IVSs. More precisely, we consider separable volatility functions of the form

$$\sigma(t,T_k) = \phi(L(t,T_k),T_k)\lambda(t,T_k), \tag{3.3}$$

$\phi : \mathbb{R}_+ \times \{T_1,\ldots,T_{n-1}\} \to \mathbb{R}_+$, which define the class of *generalised extended market models (GEMMs).* Although being somewhat more restrictive than formulation (3.2), GEMMs – as already pointed out – offer a high degree of flexibility when it comes to fitting given IVSs, while, by their separable structure, giving the user control over the volatility term-structure, and offering a high degree of numerical tractability. The price one has to pay for the added flexibility compared to the LMM- and EMM-classes is that

[1] Compare Brigo & Mercurio [2003], who also find a completely smile-consistent volatility function in this set-up. However, as opposed to our model, which will turn out to be explicit and numerically well tractable, their formulation is rather implicit in nature and hard to handle numerically.

closed formulae for the prices of standard derivatives such as caplets can no longer be expected to exist, which is in contrast to the standard LMM, and the CEV and quadratic volatility models of the EMM class. But, what is a closed-form solution anyway? As Zühlsdorff [2002] points out, the evaluation of the noncentral chi-square distributions in the CEV-caplet formula is computationally expensive and difficult to implement, and also the evaluation of the caplet-formula in the quadratic volatility setting involves a high computational effort. In sharp contrast, the analytical approximation proposed by Hagan & Woodward [1999] can be implemented and evaluated very efficiently, while at the same time being highly accurate. So even in the presence of closed caplet-price formulae, one might be better off using the Hagan-Woodward proxy outlined in the previous chapter. The bottom line of the above is that it makes little sense to artificially confine the class of functions admissible for ϕ, as this only restricts a model's ability to generate realistic IVSs, while offering little or no computational advantages. The only arguments that can be brought forward in favour of restricting oneself to a low-dimensional parametric class for ϕ are avoiding overfitting and simplifying the calibration process, but still these points are not convincing: overfitting can be ruled out by parametric classes that are not as inflexible as the above, and calibration can also be efficiently tackled for high-dimensional parametric classes, as will become clear in the following. But before we touch on these points, we will develop the general theory of GEMMs.

3.2.1 LIBOR dynamics under the forward LIBOR measure

Our next step will be to derive the dynamics of $L(\cdot, T_{k-1})$ under the $\mathbb{Q}^{T_i}$ forward-measure, with $T_i \in \mathcal{T}$, assuming throughout that σ is of the form (3.3). In the course of the derivation, we will exploit the following version of Girsanov's theorem (cf. e.g. Hunt & Kennedy [2000], p.103):

Theorem 3.1 (Girsanov). *Suppose W is a d-dimensional $(\mathbb{F}, \mathbb{P})$ Brownian motion, $\mathbb{Q} \sim \mathbb{P}$, and the strictly positive $\mathbb{P}$-martingale ζ with $\zeta(t) = \frac{d\mathbb{Q}}{d\mathbb{P}}\Big|_{\mathcal{F}_t}$ has a continuous version. Then Z defined by*

$$Z_i(t) = W_i(t) - \left[W_i, \int_0^{\cdot} \zeta(s)^{-1}\, d\zeta(s)\right](t), \quad i = 1, \ldots, d,$$

is a d-dimensional $(\mathbb{F}, \mathbb{Q})$ Brownian motion.

This leads us to the main result of this section.

Theorem 3.2 (LIBOR-dynamics under various forward measures). *Take $T_k \in \{T_1, \ldots, T_n\}$ as fixed and assume*

$$dL(t, T_{k-1}) = \sigma(t, T_{k-1}) dW^{T_k}(t),$$

where W^{T_k} is a standard d-dimensional $(\mathbb{F}, \mathbb{Q}^{T_k})$ Brownian motion. Then the following relations for the LIBOR dynamics under the forward measure $\mathbb{Q}^{T_i}$, $T_i \in \mathcal{T}$, hold:

$$i < k : dL(t, T_{k-1}) = \sigma(t, T_{k-1}) \cdot \left(\sum_{j=i+1}^{k} \frac{\tau_j}{1 + \tau_j L(t, T_{j-1})} \sigma(t, T_{j-1})' dt + dW^{T_i}(t) \right),$$

$$i > k : dL(t, T_{k-1}) = \sigma(t, T_{k-1}) \cdot \left(- \sum_{j=k+1}^{i} \frac{\tau_j}{1 + \tau_j L(t, T_{j-1})} \sigma(t, T_{j-1})' dt + dW^{T_i}(t) \right),$$

where $0 \leq t \leq \min\{T_i, T_{k-1}\}$ and W^{T_i} is a standard d-dimensional $(\mathbb{F}, \mathbb{Q}^{T_i})$-Brownian.

Proof: First we consider the case $i < k$. We may assume that $k \geq 2$, because $k = 1$ leads to $i = 0$ and $0 \leq t \leq T_0 = 0$, which is trivial. We start by deriving the LIBOR dynamics under $\mathbb{Q}^{T_{k-1}}$. In order to obtain the Radon-Nikodým derivative of $\mathbb{Q}^{T_{k-1}}$ with respect to $\mathbb{Q}^{T_k}$, we employ the change-of-numéraire technique (see e.g. Geman *et al.* [1995] or Bingham & Kiesel [2004], Chapter 9):

$$\begin{aligned} \frac{d\mathbb{Q}^{T_{k-1}}}{d\mathbb{Q}^{T_k}} &= \frac{p(0, T_k) p(T_{k-1}, T_{k-1})}{p(T_{k-1}, T_k) p(0, T_{k-1})} \\ &= \frac{1 + \tau_k L(T_{k-1}, T_{k-1})}{1 + \tau_k L(0, T_{k-1})}. \end{aligned}$$

Now for $0 \leq t \leq T_{k-1}$,

$$\begin{aligned}\left.\frac{d\mathbb{Q}^{T_{k-1}}}{d\mathbb{Q}^{T_k}}\right|_{\mathcal{F}_t} &= \mathbb{E}_{\mathbb{Q}^{T_k}}\left(\frac{d\mathbb{Q}^{T_{k-1}}}{d\mathbb{Q}^{T_k}}\middle|\mathcal{F}_t\right)\\ &= \mathbb{E}_{\mathbb{Q}^{T_k}}\left(\frac{1+\tau_k L(T_{k-1},T_{k-1})}{1+\tau_k L(0,T_{k-1})}\middle|\mathcal{F}_t\right)\\ &= \frac{1+\tau_k L(t,T_{k-1})}{1+\tau_k L(0,T_{k-1})},\end{aligned}$$

where the last equation follows because $L(\cdot,T_{k-1})$ is a $\mathbb{Q}^{T_k}$-martingale. The above version of Girsanov's theorem tells us that the process $W^{T_{k-1}}$ defined by

$$W_i^{T_{k-1}} = W_i^{T_k} - \left[W_i^{T_k}, X\right], \quad i = 1,\ldots,d,$$

with

$$\begin{aligned}X(t) &= \int_0^t \frac{1+\tau_k L(0,T_{k-1})}{1+\tau_k L(s,T_{k-1})}\, d\left(\frac{1+\tau_k L(s,T_{k-1})}{1+\tau_k L(0,T_{k-1})}\right)\\ &= \int_0^t \frac{1}{1+\tau_k L(s,T_{k-1})}\, d\left(1+\tau_k L(s,T_{k-1})\right)\\ &= \int_0^t \frac{\tau_k}{1+\tau_k L(s,T_{k-1})}\sigma(s,T_{k-1})\, dW^{T_k}(s)\end{aligned}$$

is a $\mathbb{Q}^{T_{k-1}}$-martingale. Simplifying, we get

$$\begin{aligned}dW_i^{T_{k-1}}(t) &= dW_i^{T_k}(t) - d\left[W_i^{T_k}, X\right](t)\\ &= dW_i^{T_k}(t) - dW_i^{T_k}(t)dX(t)\\ &= dW_i^{T_k}(t) - dW_i^{T_k}(t)\frac{\tau_k}{1+\tau_k L(t,T_{k-1})}\sigma(t,T_{k-1})dW^{T_k}(t),\end{aligned}$$

or in vector form

$$\begin{aligned}dW^{T_{k-1}}(t) &= dW^{T_k}(t) - \frac{\tau_k}{1+\tau_k L(t,T_{k-1})}dW^{T_k}(t)\,\sigma(t,T_{k-1})\,dW^{T_k}(t)\\ &= dW^{T_k}(t) - \frac{\tau_k}{1+\tau_k L(t,T_{k-1})}dW^{T_k}(t)\,dW^{T_k}(t)'\,\sigma(t,T_{k-1})'\\ &= dW^{T_k}(t) - \frac{\tau_k}{1+\tau_k L(t,T_{k-1})}\sigma(t,T_{k-1})'\,dt.\end{aligned}$$

For the LIBOR dynamics under $\mathbb{Q}^{T_{k-1}}$ we obtain

$$\begin{aligned}dL(t,T_{k-1}) &= \sigma(t,T_{k-1})dW^{T_k}(t)\\ &= \sigma(t,T_{k-1})\left(\frac{\tau_k}{1+\tau_k L(t,T_{k-1})}\sigma(t,T_{k-1})'dt + dW^{T_{k-1}}(t)\right),\end{aligned}$$

which completes the proof for the case $i = k - 1$. The corresponding result for general $i < k - 1$ can now easily be obtained via backward induction by repeating the above argument. The case $i > k$ can be handled in a similar fashion. □

An immediate consequence of the above theorem is that a forward LIBOR process $L(\cdot, T_{k-1})$ is a martingale only under its respective forward measure $\mathbb{Q}^{T_k}$. In the standard LIBOR market model, the above result also implies that a forward LIBOR rate is log-normal only under its respective forward measure.

The following existence theorem for GEMMs slightly generalises the corresponding result for EMMs due to Andersen & Andreasen [2000].

Theorem 3.3. *Suppose that* $L(0, T_{k-1}) \geq 0$ *for all* $k \in \{1, \ldots, n\}$. *If* ϕ

(i) is locally Lipschitz continuous, i.e. $\forall z \; \exists c(z) > 0$ *such that for* $0 \leq x, y < z$ *and* $\forall k \in \{0, \ldots, n-1\} : |\phi(x, T_k) - \phi(y, T_k)| \leq c(z)|x - y|$, *and*

(ii) satisfies a linear growth condition:

$$\exists c > 0 \text{ such that } \forall x \geq 0 \text{ and } \forall k \in \{0, \ldots, n-1\} : \phi(x, T_k)^2 \leq c(1 + x)^2,$$

then a non-explosive, unique solution of the system of SDEs in Theorem 3.2 exists under all measures $\mathbb{Q}^{T_k}$. *If further* $L(0, T_{k-1}) > 0$ *for all* $k \in \{1, \ldots, n\}$, *then the solution is positive for all* $t \geq 0$.

The local Lipschitz condition in the above theorem guarantees the uniqueness of the solution, the linear growth condition its non-explosiveness in finite time. The proof, which uses standard existence and uniqueness arguments, proceeds exactly along the same lines as the one in Andersen & Andreasen [2000], to which we refer for further details.

In general, it is not possible to state an explicit solution of our system of SDEs, not even for the standard LIBOR market model. As a consequence, one has to resort to numerical methods such as Monte Carlo simulation when pricing certain complex derivatives that depend on the simultaneous realisation of several LIBOR rates in the above setup. The relation given in the following corollary is central for the simulation of the LIBOR market

model, as for simulation purposes *one* measure $\mathbb{Q}^{T_k}$ has to be chosen, under which all forward LIBOR rates have to be evolved simultaneously.

Corollary 3.1. *Under the assumptions of Theorem 3.2, we find the following relation between the Brownian motions W^{T_k} and $W^{T_{k-1}}$ under the respective measures $\mathbb{Q}^{T_k}$ and $\mathbb{Q}^{T_{k-1}}$:*

$$dW^{T_{k-1}}(t) = dW^{T_k}(t) - \frac{\tau_k}{1+\tau_k L(t,T_{k-1})}\sigma(t,T_k)'\,dt.$$

Supposing that we choose e.g. the *terminal measure* $\mathbb{Q}^{T_n}$ as a reference, the above corollary helps us to inductively construct the LIBOR processes $L(\cdot,T_{n-1})$, $L(\cdot,T_{n-2})$ etc. under $\mathbb{Q}^{T_n}$.

3.3 Pricing caplets and floorlets in a GEMM

A *caplet with reset date T_k, maturity T_{k+1} and strike rate K*, or briefly a *T_k-caplet with strike K*, is a derivative that pays the holder

$$\tau_{k+1}(L(T_k,T_k)-K)^+$$

at time T_{k+1}. Hence a caplet can be regarded as a call option on a LIBOR rate, or, equivalently, an insurance against interest rates rising above a certain level. A *T_k-floorlet with strike K* is a derivative that pays the holder

$$\tau_{k+1}(K-L(T_k,T_k))^+$$

at time T_{k+1}. The arbitrage-free price of a T_k-caplet at time t is

$$p(t,T_{k+1})\mathbb{E}_{\mathbb{Q}^{T_{k+1}}}\left[\frac{\tau_{k+1}(L(T_k,T_k)-K)^+}{p(T_{k+1},T_{k+1})}\middle|\mathcal{F}_t\right]$$
$$= \tau_{k+1}p(t,T_{k+1})\mathbb{E}_{\mathbb{Q}^{T_{k+1}}}\left[(L(T_k,T_k)-K)^+\middle|\mathcal{F}_t\right]$$

The Feynman-Kac connection now allows us to state the following theorem.

Theorem 3.4. *The t-price $C(t,T_k,K)$ of a T_k-caplet with strike K is given by*

$$C(t,T_k,K) = \tau_{k+1}p(t,T_{k+1})f(L(t,T_k),v(t,T_k))$$

with

$$v(t,T_k) = \int_t^{T_k} \|\lambda(s,T_k)\|^2 \, ds,$$

where $f(x,\tau)$ *solves the initial value problem*

$$\frac{1}{2}\phi(x,T_k)^2 \frac{\partial^2}{\partial x^2} f(x,\tau) = \frac{\partial}{\partial \tau} f(x,\tau)$$

with initial condition

$$f(x,0) = (x-K,0)^+.$$

A similar result holds for floorlets, with the difference that the initial condition is $f(x,0) = (K-x,0)^+$. The proof, which uses a deterministic time-change argument and the Feynman-Kac formula, can be found in Andersen & Andreasen [2000]. The above PDE can be solved efficiently by numerical methods, e.g. a Crank-Nicolson finite difference scheme. In addition, it is also possible to use the Hagan-Woodward proxy to get approximate caplet-prices. As *caps* and *floors* are portfolios of caplets and floorlets, the above results readily extend to the pricing of these instruments.

3.4 Pricing swaptions in a GEMM

An interest rate swap (IRS) is a contract to exchange fixed against floating payments, where the floating payments typically depend on LIBOR rates. An IRS is specified by its *reset-dates* $T_\alpha, T_{\alpha+1}, \ldots, T_{\beta-1}$, its *payment-dates* $T_{\alpha+1}, \ldots, T_\beta$, and the fixed rate K. At every $T_j \in \{T_{\alpha+1}, \ldots, T_\beta\}$, the fixed payment is $\tau_j K$ with $\tau_j = T_j - T_{j-1}$, while the floating payment is $\tau_j L(T_{j-1}, T_{j-1})$. That is, the floating payment in T_j is already determined in T_{j-1}. The set of fixed (floating) payments is called the *fixed (floating) leg* of the swap, and the party that makes the fixed (floating) payments is said to hold a *payer (receiver) swap*. The value of a swap in $t \leq T_\alpha$ can be determined without making any distributional assumptions on the LIBOR rates, as the following considerations show. The T_{j-1}-value of the floating payment $\tau_j L(T_{j-1}, T_{j-1})$ that is paid (respectively received) at

time T_j is

$$p(T_{j-1},T_j)\tau_j L(T_{j-1},T_{j-1}) = p(T_{j-1},T_j)\tau_j \frac{1}{\tau_j}\left(\frac{1}{p(T_{j-1},T_j)} - 1\right)$$
$$= 1 - p(T_{j-1},T_j),$$

and therefore the t-value $(t \leq T_{j-1})$ of the floating payment must be $p(t,T_{j-1}) - p(t,T_j)$. It follows that the value of a payer swap in $t \leq T_\alpha$ is

$$\sum_{i=\alpha+1}^{\beta} (p(t,T_{i-1}) - p(t,T_i)) - \sum_{i=\alpha+1}^{\beta} p(t,T_i)\tau_i K$$
$$= \sum_{i=\alpha+1}^{\beta} (p(t,T_{i-1}) - (1+\tau_i K)p(t,T_i))$$
$$= \sum_{i=\alpha+1}^{\beta} p(t,T_i)\tau_i (L(t,T_{i-1}) - K), \tag{3.4}$$

or alternatively

$$\sum_{i=\alpha+1}^{\beta} (p(t,T_{i-1}) - p(t,T_i)) - \sum_{i=\alpha+1}^{\beta} p(t,T_i)\tau_i K$$
$$= p(t,T_\alpha) - p(t,T_\beta) - \sum_{i=\alpha+1}^{\beta} p(t,T_i)\tau_i K, \tag{3.5}$$

since other prices obviously give rise to arbitrage opportunities. Accordingly, the value of a receiver swap is

$$\sum_{i=\alpha+1}^{\beta} p(t,T_i)\tau_i (K - L(t,T_{i-1}))$$
$$= p(t,T_\beta) - p(t,T_\alpha) + \sum_{i=\alpha+1}^{\beta} p(t,T_i)\tau_i K.$$

The *forward swap rate* (FSR) at time t of the above IRS, which we denote by $S_{\alpha,\beta}(t)$, is the value for the fixed rate K that makes the t-value of the IRS zero. $S_{\alpha,\beta}(t)$ can thus be obtained by equating expression (3.4) to zero and solving for K, which gives

$$S_{\alpha,\beta}(t) = \sum_{i=\alpha+1}^{\beta} w_i(t) L(t,T_{i-1}) \tag{3.6}$$

with

$$w_i(t) = \frac{\tau_i p(t,T_i)}{\sum_{j=\alpha+1}^{\beta} \tau_j p(t,T_j)},$$

or equivalently by equating (3.5) to zero, which gives

$$S_{\alpha,\beta}(t) = \frac{p(t,T_\alpha) - p(t,T_\beta)}{\sum_{i=\alpha+1}^{\beta} \tau_i p(t,T_i)}.$$

Equation (3.6) shows that the FSR can be expressed as a suitably weighted average of the spanning forward LIBORs.

Now we introduce *swap options*, or *swaptions* for short, which, along with caps and floors, constitute the most popular instruments in the interest rate derivatives world. A European *payer swaption* gives the holder the right to enter a swap as fixed-rate payer at a fixed rate K (the *swaption strike*) at a future date that coincides with the first reset date T_α of the underlying swap. Similarly, a European *receiver swaption* gives the holder the right to enter a swap as fixed-rate receiver. To express the value of a payer swaption as a function of the swap rate, first notice that the following relation, which can easily be verified, holds for the t-value of a payer swap: For $0 \leq t \leq T_\alpha$,

$$\sum_{i=\alpha+1}^{\beta} p(t,T_i)\tau_i(L(t,T_{i-1}) - K) = (S_{\alpha,\beta}(t) - K) \sum_{i=\alpha+1}^{\beta} \tau_i p(t,T_i). \tag{3.7}$$

Needless to say, a similar formula can be derived for receiver swaps. The advantage of the expression on the right-hand side of equation (3.7) over the expression on the left-hand side (which is our formula (3.4)) is that one can instantly tell from $S_{\alpha,\beta}(t)$ if the t-value of the payer swap is positive or negative, which is not at all obvious from formula (3.4). At the maturity date T_α, a payer swaption is exercised if and only if the value of the underlying swap is positive, which is the case if and only if $S_{\alpha,\beta}(T_\alpha) - K > 0$ holds. Clearly, the payer swaption-value in T_α is

$$(S_{\alpha,\beta}(T_\alpha) - K)^+ \sum_{i=\alpha+1}^{\beta} \tau_i p(T_\alpha,T_i),$$

and the receiver swaption-value is

$$(K - S_{\alpha,\beta}(T_\alpha))^+ \sum_{i=\alpha+1}^{\beta} \tau_i p(T_\alpha,T_i).$$

3.4.1 An approximate pricing formula for swaptions

Now we turn to the problem of deriving an approximate pricing formula for swaptions. In an LMM, swaptions generally have to be priced by Monte Carlo simulation. For calibration purposes, it can be helpful to have an analytical approximation at hand, as this can considerably speed up the calibration process. Analytical approximations can also be useful for calculating sensitivities, because sensitivities obtained by Monte Carlo simulation are notoriously unstable and prone to numerical errors.[2] We observe that

$$A_{\alpha,\beta}(t) = \sum_{i=\alpha+1}^{\beta} \tau_i p(t, T_i)$$

is the t-price of a portfolio of bonds (i.e. a traded asset). Consequently, $A_{\alpha,\beta}(t)$, which if known as *accrual factor* or *present value of a basis point,* can be used as numéraire. Now note that

$$S_{\alpha,\beta}(t) = \frac{p(t, T_\alpha) - p(t, T_\beta)}{A_{\alpha,\beta}(t)},$$

where the numerator can be regarded as the price of a traded asset as well. We conclude that, in order for our model to be arbitrage-free, the swap rate $S_{\alpha,\beta}$ has to be a martingale under the numéraire pair $(\mathbb{Q}^{\alpha,\beta}, A_{\alpha,\beta}(\cdot))$. $\mathbb{Q}^{\alpha,\beta}$ is the so-called *forward swap measure.*

In order to derive a tractable SDE for $(S_{\alpha,\beta}(t))$ given by (3.6), we need some simplifying assumptions. First, for most reasonable shifts of the LIBOR curve, the weights $w_i(t)$ vary only little, and therefore can be approximated by their initial values $w_i(0)$ ('freezing the weights').[3] Second, if the LIBOR curve experiences predominantly parallel shifts (as is the case in practice), approximating

$$\frac{\phi(L(t, T_i), T_i)}{\phi_S(S_{\alpha,\beta}(t))}$$

with $\phi_S(\cdot)$ defined by

$$\phi_S(S_{\alpha,\beta}(t)) = \sum_{i=\alpha+1}^{\beta} w_i(0)\phi(S_{\alpha,\beta}(t), T_{i-1})$$

[2] See also Glasserman & Zhao [1999], who deal with the problem of calculating sensitivities in an LMM by Monte Carlo methods, and present algorithms that substantially improve on a naïve Monte Carlo approach in terms of quality and speed.

[3] This can be shown by simulation studies.

by its initial value $\frac{\phi(L(0,T_i),T_i)}{\phi_S(S_{\alpha,\beta}(0))}$ is often reasonable. Third, assuming zero drifts of the LIBOR processes under the forward swap-measure will typically have a negligible impact on swaption prices ('collapsing all measures'). Applying these simplifications leads to

$$\begin{aligned}
dS_{\alpha,\beta}(t) &\approx \sum_{i=\alpha+1}^{\beta} w_i(0)\, dL(t,T_{i-1}) \\
&\approx \sum_{i=\alpha+1}^{\beta} w_i(0)\phi(L(t,T_{i-1}),T_{i-1})\lambda(t,T_{i-1})\, dW_{\alpha,\beta}(t) \\
&= \sum_{i=\alpha+1}^{\beta} w_i(0)\frac{\phi(L(t,T_{i-1}),T_{i-1})}{\phi_S(S_{\alpha,\beta}(t))}\phi_S(S_{\alpha,\beta}(t))\lambda(t,T_{i-1})\, dW_{\alpha,\beta}(t) \\
&\approx \sum_{i=\alpha+1}^{\beta} w_i(0)\frac{\phi(L(0,T_{i-1}),T_{i-1})}{\phi_S(S_{\alpha,\beta}(0))}\phi_S(S_{\alpha,\beta}(t))\lambda(t,T_{i-1})\, dW_{\alpha,\beta}(t) \\
&= \phi_S(S_{\alpha,\beta}(t)) \sum_{i=\alpha+1}^{\beta} \omega_i\lambda(t,T_{i-1})\, dW_{\alpha,\beta}(t) \qquad (3.8)
\end{aligned}$$

with

$$\omega_i = w_i(0)\frac{\phi(L(0,T_{i-1}),T_{i-1})}{\phi_S(S_{\alpha,\beta}(0))}$$

and $W_{\alpha,\beta}(\cdot)$ a standard d-dimensional Brownian motion under $\mathbb{Q}^{\alpha,\beta}$.

For the price $S(0,T_\alpha,T_\beta,K)$ of a payer swaption with strike K at time $t=0$, we get

$$\frac{S(0,T_\alpha,T_\beta,K)}{A_{\alpha,\beta}(0)} = \mathbb{E}_{\mathbb{Q}^{\alpha,\beta}}\left[\frac{(S_{\alpha,\beta}(T_\alpha)-K)^+}{A_{\alpha,\beta}(T_\alpha)}A_{\alpha,\beta}(T_\alpha)\right],$$

which reduces to

$$S(0,T_\alpha,T_\beta,K) = A_{\alpha,\beta}(0)\mathbb{E}_{\mathbb{Q}^{\alpha,\beta}}\left[(S_{\alpha,\beta}(T_\alpha)-K)^+\right].$$

Approximating the swap rate process by (3.8) leads us to conclude that we can calculate the approximate price of above swaption by the same method as caplet-prices in the GEMM:

Theorem 3.5. *The price $S(0,T_\alpha,T_\beta,K)$ of a payer swaption in a GEMM can be approximated by*

$$S(0,T_\alpha,T_\beta,K) \approx A_{\alpha,\beta}(0)g(S_{\alpha,\beta}(0),v_S(0,T_\alpha)) \qquad (3.9)$$

with

$$v_S(0,T_\alpha) = \int_0^{T_\alpha} \left\| \sum_{i=\alpha}^{\beta-1} \omega_i \lambda(s,T_i) \right\|^2 ds,$$

where $g(x,\tau)$ *solves the initial value problem*

$$\frac{1}{2}\phi_S(x)^2 \frac{\partial^2}{\partial x^2} g(x,\tau) = \frac{\partial}{\partial \tau} g(x,\tau)$$

with initial condition

$$g(x,0) = (x-K,0)^+.$$

The above price can be worked out numerically, or, as for caplets, the Hagan-Woodward formula can be used to get an approximation for (3.9). Hull & White [2000] derive a somewhat more exact approximation in the standard LMM by proceeding along similar lines. Their approximation can also be readily extended to our GEMM setting.

3.4.2 Numerical tests

The quality of the proposed approximation will generally depend on the concrete choice of functions and parameters involved in the (exact) swap rate dynamics. While it is not feasible to test the approximation for all realistic scenarios, a closer look at a few typical ones might already suffice to convey a good impression of its quality. We choose to study the scenarios detailed in Table 3.1.[4]

We assume that the current time t is 0, and calculate the prices of swaptions with maturities of 1, 5 and 10 years on swaps that run for 1, 5 and 10 years and are reset semi-annually, i.e. $\tau_i = \tau = 0.5$ for all i.[5] These test-cases are inspired by those in Andersen & Andreasen [2000]. The choice of closely related test-scenarios enables us to compare our results with those of Andersen and Andreasen, whereby we can gauge the impact of allowing for a differential volatility structure (with $\phi(\cdot,T_i)$ depending on T_i rather than using one fixed ϕ for all maturities) on the quality of the approximation. For all three scenarios, the motivation behind the choice of the local volatility function ϕ is the empirically observed flattening-out of the IVS with increasing times to maturity. Scenario 1 corresponds

[4] In the table, $[T-t]$ denotes the integer part of $T-t$.

[5] For example, a swaption maturing in 1 year on a swap running for 5 years is called a 1 into 5 swaption.

Scenario 1
$L(0,T_i) = 0.06,\ T_i \in \{0, 0.5, \ldots, 19.5\}$
$\lambda(t,T_i) \equiv 0.05,\ T_i \in \{0.5, 1, \ldots, 19.5\}$
$\phi(x,T_i) = x \min\left\{x^{0.05T_i - 1}, 20\right\},\ T_i \in \{0.5, 1, \ldots, 19.5\}$
Scenario 2
$L(0,T_i) = 0.06,\ T_i \in \{0, 0.5, \ldots, 19.5\}$
$\lambda(t,T_i) = \max\{0.1 - 0.01[T_i - t], 0.01\},\ T_i \in \{0.5, 1, \ldots, 19.5\}$
$\phi(x,T_i) = x \min\left\{x^{0.025T_i - 0.5}, 20\right\},\ T_i \in \{0.5, 1, \ldots, 19.5\}$
Scenario 3
$L(0,T_i) = 0.06,\ T_i \in \{0, 0.5, \ldots, 19.5\}$
$\lambda(t,T_i) = (\lambda_1(t,T_i), \lambda_2(t,T_i)) = \left(0.04, 0.04 - 0.015\sqrt{[T_i - t]}\right),\ T_i \in \{0.5, 1, \ldots, 19.5\}$
$\phi(x,T_i) = x \min\left\{x^{0.025T_i - 0.5}, 20\right\},\ T_i \in \{0.5, 1, \ldots, 19.5\}$

Table 3.1: Test scenarios

to a simple one-factor model with constant λ. Scenario 2 builds on a decreasing function λ which takes account of the empirically motivated rule of thumb that the volatility of LIBOR rates decreases with increasing time to maturity. Finally, scenario 3 is a two-factor model, where λ_1 represents parallel shifts of the term structure, and λ_2 changes of its steepness (often called twists).[6] In all scenarios, λ is chosen to be stationary, in the sense that it only depends on the time to maturity $T - t$, which is desirable from an empirical viewpoint (see Brigo & Mercurio [2001]). Figures 3.1 and 3.2 show $\phi(x,T)/x$ and $\lambda(t,T)$ for scenarios 2 and 3, respectively.[7]

We pursue the following testing plan: For every swaption, we perform a Monte Carlo simulation with 10 million antithetic paths. To keep the discretization-bias arising in the course of the simulation-procedure as low as possible, we use a step-size of 0.0625 (years) for the simulation, and choose piecewise constant functions λ.[8] Then we calculate the means and standard errors reported in Tables 3.2, 3.3 and 3.4. The approximations in the

[6] This choice can be justified by principal component analysis of empirical term structure data.

[7] For comparison: In a standard LMM (cf. (3.1)), we have $\phi(x,T)/x \equiv 1$.

[8] For practical applications, one would typically choose a significantly larger step-size to speed up the simulation procedure without losing much accuracy.

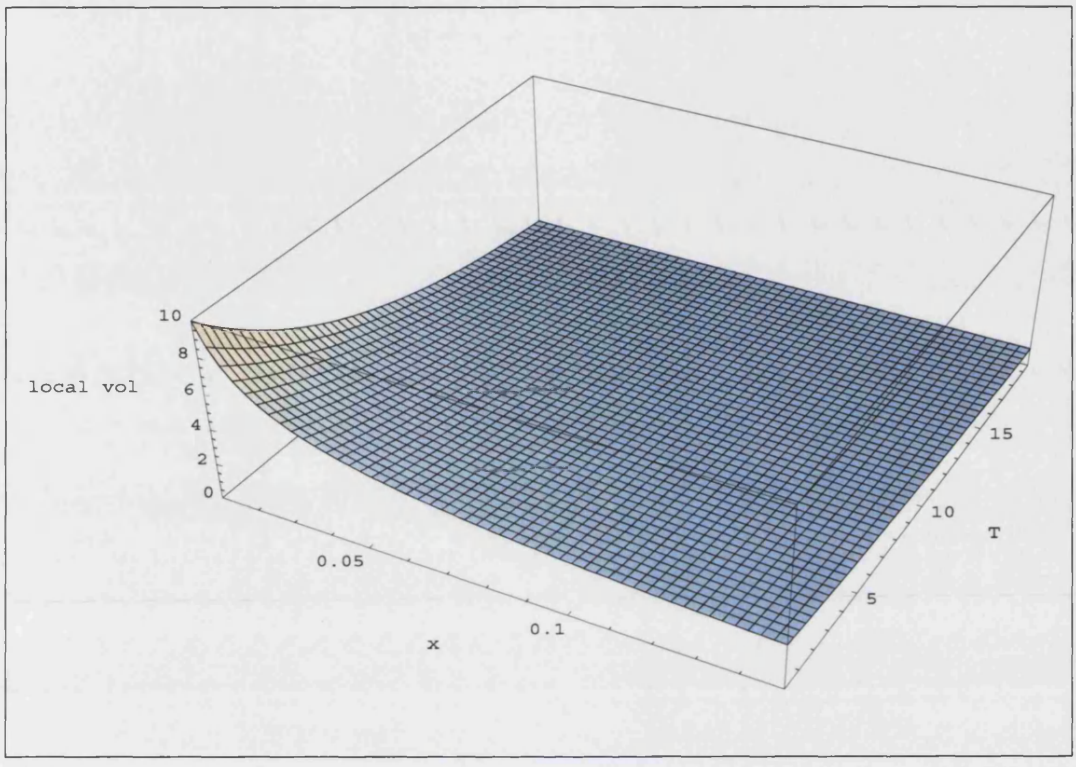

Figure 3.1: $\phi(x,T)/x$ for scenarios 2 and 3

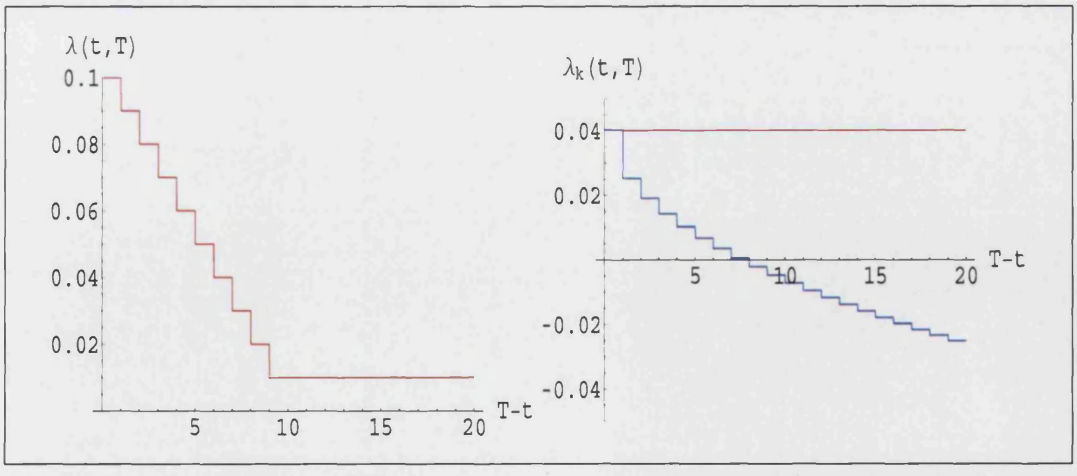

Figure 3.2: $\lambda(t,T)$ for scenarios 2 and 3

Swaption	T_α	T_β	Simul. price	Std. err.	Proxy	Abs. error	Rel. err. in %
1×1	1	2	147.797	0.102	146.612	-1.186	-0.802
1×5	1	6	520.820	0.579	522.522	1.702	0.327
1×10	1	11	712.767	0.706	717.940	5.172	0.726
5×1	5	6	148.672	0.103	148.517	-0.156	-0.105
5×5	5	10	520.749	0.252	523.227	2.477	0.476
5×10	5	15	712.131	0.400	718.672	6.541	0.918
10×1	10	11	78.099	0.072	78.055	-0.044	-0.056
10×5	10	15	273.027	0.200	273.567	0.540	0.198
10×10	10	20	373.457	0.201	375.406	1.949	0.522

Table 3.2: Simulated and approximate prices of at-the-money payer swaptions for scenario 1 in basis points

tables are obtained by numerically solving (3.9) with a Crank-Nicolson finite difference scheme.

Even a superficial glance at the tables reveals that the quality of approximation in our test cases is excellent, the approximation error never exceeding 1%. A comparison with the approximation errors reported by Andersen and Andreasen for their test-cases shows that they are of the same order of magnitude, and – at least in the scenarios under consideration – our more general approximation that allows for maturity-dependent functions ϕ does about equally well as their limited one. As already mentioned, testing could be extended to a sufficiently rich set of scenarios, in order to get a better impression of the reliability of our approximation and the determinants of its quality. For reasons of scope, we will refrain from doing so at this point, and leave this for further research.

3.5 GEMMs in practice

Numerical issues are of paramount importance for the practical applicability of a financial model, even more so if it is geared to become a uniformly accepted market standard. In the following, we will briefly touch on questions concerning calibration and simulation in

Swaption	T_α	T_β	Simul. price	Std. err.	Proxy	Abs. error	Rel. err. in %
1 × 1	1	2	78.506	0.102	78.465	-0.041	-0.053
1 × 5	1	6	253.113	0.313	253.456	0.343	0.135
1 × 10	1	11	303.541	0.309	304.126	0.585	0.193
5 × 1	5	6	84.064	0.098	84.109	0.044	0.053
5 × 5	5	10	258.899	0.324	259.898	0.999	0.386
5 × 10	5	15	301.658	0.283	303.441	1.783	0.591
10 × 1	10	11	47.495	0.064	47.509	0.015	0.031
10 × 5	10	15	142.783	0.070	143.066	0.283	0.198
10 × 10	10	20	167.378	0.168	168.006	0.628	0.375

Table 3.3: Simulated and approximate prices of at-the-money payer swaptions for scenario 2 in basis points

Swaption	T_α	T_β	Simul. price	Std. err.	Proxy	Abs. error	Rel. err. in %
1 × 1	1	2	43.714	0.046	43.863	0.149	0.072
1 × 5	1	6	144.073	0.163	144.051	-0.022	-0.015
1 × 10	1	11	205.107	0.212	204.758	-0.349	-0.004
5 × 1	5	6	49.071	0.050	49.045	-0.026	-0.052
5 × 5	5	10	177.357	0.201	177.423	0.066	0.037
5 × 10	5	15	261.717	0.181	262.017	0.300	0.115
10 × 1	10	11	34.200	0.024	34.188	-0.012	-0.034
10 × 5	10	15	128.764	0.059	128.743	-0.021	-0.016
10 × 10	10	20	194.403	0.194	194.512	0.109	0.056

Table 3.4: Simulated and approximate prices of at-the-money payer swaptions for scenario 3 in basis points

the GEMM-context. For an exhaustive treatment of these points in a standard LMM context, we refer to Brace *et al.* [1998], Glasserman & Zhao [1999], Glasserman & Zhao [2000], Hull & White [2000], Pelsser [2000], Brigo & Mercurio [2001], Rebonato [2002], Rebonato [2003] and Rebonato [2004].

3.5.1 Calibration

The only major difference between the standard LMM and the GEMM with regard to calibration relates to the choice of the functions $\phi(\cdot, T_k)$: Once the IVSs for all maturities T_k have been obtained (e.g. by fitting functions of a certain parametric class to market-observed implied volatilities), one of the approximations (P1) to (P6) is chosen to calculate the $\phi(\cdot, T_k)$; this already ensures smile-consistency, without having to go through computationally expensive and often unstable calibration procedures. After this step, one can proceed in much the same way as for standard LMMs. For example, the functions $\lambda(\cdot, T_k)$ can be chosen so as to reflect a trader's view on future IVSs, and the remaining degrees of freedom (e.g. the number of driving factors) can be utilised to calibrate the model to swaption prices and/or historical LIBOR correlations.

3.5.2 Simulation

The simulation procedures for EMMs (as outlined e.g. in Andersen & Andreasen [2000]) can be adapted in a straightforward fashion to the GEMM-setting. Perhaps a bit surprisingly, despite the complex LIBOR-dynamics, GEMM-simulations can be carried out as efficiently as LMM-simulations, provided that the values of the functions $\phi(\cdot, T_k)$ are tabulated before running the simulation. Thereby, the evaluation of ϕ after each simulation-step can be avoided. The numerical error introduced by tabulating the $\phi(\cdot, T_k)$, which amounts to approximating them by piecewise constant functions, is negligible as long as the discretization-mesh is reasonably fine. Another issue relates to the choice of the step-size in the simulation process; this question was explored by Andersen & Andreasen [2000] in the EMM setting. At this point, suffice it to say that simulation studies we conducted for different classes of ϕ have shown that using the same step-sizes as the above-mentioned authors for the EMM also leads to negligible discretization errors in the GEMM.

3.6 Summary

In this section, we develop the theory of generalised extended market models. Building on and extending results from Andersen & Andreasen [2000], we show that there exist efficient approximations for both caps/floors and swaptions, regardless of the specific form of ϕ. The calibration procedures for standard LIBOR market models that rely on efficient approximations can therefore be used in the GEMM setting. Combined with our results from the previous chapter, this means that we are able to fit any given set of implied volatility smiles for the maturities of the tenor structure (almost) exactly, while at the same time preserving numerical tractability. Calibration to caplet smiles is, by the methods outlined in the previous chapter, almost immediate, and calibration to volatility term structures and swaption prices can follow the same methods as in the standard LMM. Summing up, we propose a very flexible framework that is both smile-consistent and numerically highly tractable.

Chapter 4

LIBOR market models driven by Lévy processes

Modelling equity prices through Lévy processes has become very popular in recent years and has been the subject of many studies (see e.g. Eberlein & Keller [1995], Barndorff-Nielsen [1998], Chan [1999], Prause [1999], Raible [2000], Barndorff-Nielsen *et al.* [2002], Carr *et al.* [2002], Carr *et al.* [2003], Bingham & Kiesel [2004]). Lévy-driven models improve on many of the shortcomings of the classical Black-Scholes approach. As pointed out in numerous publications, empirical evidence strongly contradicts the log-normality and path-continuity assumptions underlying the Black-Scholes model: Empirical log-return distributions are typically skewed and more leptokurtic than the normal distribution, and jumps in price processes can easily be identified.[1] Lévy processes, in contrast, are much better able to capture these stylised properties. What makes Lévy processes of particular importance for our purposes is the inherent deviation from the log-normality assumption, in combination with the great flexibility they provide when it comes to fitting marginal return distributions, as these features allow for an excellent fit to observed implied volatility surfaces (see e.g. Schoutens *et al.* [2003] for a treatment in the equity context).

However, research in interest rate models with Lévy processes as driving noise is still in its

[1] The simple fact alone that trading takes place only at discrete time-points causes discontinuities in price processes.

infancy. The relative sparseness of research in this area seems to be due to various reasons. Firstly, the development in mathematical finance during the last decades clearly shows that extensions and modifications of the classical Black-Scholes paradigm, such as incorporating jumps, local or stochastic volatility in price processes, have first been extensively studied in the equity or foreign exchange context before finding their way into the interest rate world. Evidently, this is due to the additional mathematical complexity one faces in the construction of interest rate models, which prevents approaches that work in the equity world from being carried over in a straightforward fashion. Secondly, only in recent years have advanced interest rate models, such as the Heath-Jarrow-Morton [1992] framework and most notably the LIBOR market models (Brace *et al.* [1996], Jamshidian [1997], Miltersen *et al.* [1997]) been thoroughly understood and embraced by both practitioners and the academic community. Now that these models are at a relatively mature stage and their practical application is under control, an extension to the Lévy context seems to be a natural step.

To the best of our knowledge, the only treatments of Lévy-driven LIBOR market models are Özkan [2002] and Eberlein & Özkan [2004].[2] However, those are purely theoretical, and the reader is left without a clear idea regarding the (non-trivial) implementation and related numerical issues. In what follows, we shall make an attempt to at least partially bridge this gap. After laying out the mathematical basis, we shall, using a change of numéraire argument, give a novel derivation of the relationships between the various forward LIBOR measures, which are the essential building blocks of our model. Then, we will proceed by proposing approximation techniques that substantially facilitate implementation. Their quality will subsequently be studied using a concrete parameterisation. In this context, we will also discuss further issues concerning simulation and implementation.

4.1 Lévy processes, additive processes and beyond

First, we give a short introduction to Lévy processes and compile some important results in order to render a self-contained treatment of Lévy-based LMMs possible. More detailed

[2]But see Glasserman & Kou [2003], Glasserman & Merener [2003a] and Glasserman & Merener [2003b] for a jump-diffusion LIBOR market model and related numerical methods and approximations.

accounts on Lévy processes can be found e.g. in Bertoin [1998], Sato [1999], Geman [2002], Protter [2003], Schoutens [2003], Applebaum [2004], Bingham & Kiesel [2004] and Cont & Tankov [2004].

4.1.1 Lévy processes

Definition 4.1. *Assume a filtered probability space* $(\Omega, \mathbb{F} = (\mathcal{F}_t), \mathcal{F}, \mathbb{P})$ *satisfying the usual conditions. A càdlàg*[3]*, adapted process* $X = (X_t)_{t \geq 0}$ *with* $X_0 = 0$ *a.s. is a* Lévy process *if*[4]

(i) X *has increments independent of the past, i.e.* $X_t - X_s$ *is independent of* $\mathcal{F}_s$, $0 \leq s \leq t < \infty$.

(ii) X *has stationary increments, i.e.* $X_t - X_s$ *has the same distribution as* X_{t-s}, $0 \leq s \leq t < \infty$.

(iii) X *is continuous in probability, i.e.* $\forall \epsilon > 0$ *and* $\forall t \geq 0$ *we have* $\lim_{h \to 0} \mathbb{P}(|X_{t+h} - X_t| > \epsilon) = 0$.

It follows immediately from the defining properties that X_t can be expressed as

$$X_t = \left(X_t - X_{t\frac{n-1}{n}}\right) + \left(X_{t\frac{n-1}{n}} - X_{t\frac{n-2}{n}}\right) + \cdots + X_{\frac{t}{n}},$$

that is as the sum of $n \in \mathbb{N}$ independent and identically distributed random variables: X_t has an *infinitely divisible* distribution. The following theorem clarifies the relation between infinitely divisible distributions and Lévy processes (see e.g. Cont & Tankov [2004], p.69):

Theorem 4.1. *Let* $X = (X_t)_{t \geq 0}$ *be a Lévy process. Then for every* t, X_t *has an infinitely divisible distribution. Conversely, if* F *is an infinitely divisible distribution, then there exists a Lévy process* X *such that* F *is the distribution of* X_1.

[3] In principle, it is possible to define a Lévy processes without imposing the càdlàg-property. But it can be shown (see e.g. Protter [2003], Chapter 2, Theorem 30), that every Lévy process (defined without the càdlàg-property) has a unique càdlàg modification. So there is no loss of generality in imposing the càdlàg-condition.

[4] For notational simplicity, we consider only $\mathbb{R}$-valued processes.

A major result, which relates the characteristic function of X_t to the characteristic triplet of the Lévy process X, is the Lévy-Khintchine representation (see e.g. Cont & Tankov [2004], p.83):

Theorem 4.2 (Lévy-Khintchine representation). *Let $X = (X_t)_{t\geq 0}$ be a Lévy process. Then the characteristic function of X_t is*

$$\mathbb{E}\left[e^{iuX_t}\right] = e^{-t\phi(u)}$$

with

$$\phi(u) = \frac{b^2}{2}u^2 - i\gamma u + \int_{|x|\geq 1}\left(1 - e^{iux}\right)\,\nu(dx) + \int_{|x|<1}\left(1 - e^{iux} + iux\right)\,\nu(dx),$$

where $\gamma \in \mathbb{R}$ is called the drift *of the Lévy process and $b^2 \geq 0$ the* diffusion coefficient. *ν is a positive measure on $\mathbb{R}$ with $\nu(\{0\}) = 0$ such that*

$$\int_{\mathbb{R}} \min\{1, x^2\}\,\nu(dx) < \infty$$

and is called the Lévy measure *of the process X, and (γ, b^2, ν) is called the* Lévy triplet *or* characteristic triplet *of the process X.*[5]

The next theorem – the Lévy-Itô decomposition – is intimately related to the Lévy-Khintchine representation. The close link between these fundamental results is detailed e.g. in Bingham & Kiesel [2004], p.183, and Cont & Tankov [2004], p.79. But first, we need the following definition (see Sato [1999], p.119).

Definition 4.2 (Poisson random measure). *Set $I = (0,\infty) \times \mathbb{R}$, and denote the Borel σ-algebra of I by $\mathcal{B}(I)$. Let $(I, \mathcal{B}(I), \nu)$ be a σ-finite measure space. A family of integer-valued, non-negative random variables $\{Q(B) : B \in \mathcal{B}(I)\}$ is called* Poisson random measure *on I with Lévy (or intensity) measure ν, if the following conditions hold:*

(i) For every $B \in \mathcal{B}(I)$, $Q(B)$ follows a Poisson-distribution with mean $\nu(B)$.

(ii) If $B_1, \ldots, B_n \in \mathcal{B}(I)$ are disjoint, then $Q(B_1), \ldots, Q(B_n)$ are independent.

[5] It is worth noting that the drift γ depends on the truncation function being used; see e.g. Cont & Tankov [2004], p.83, or Bingham & Kiesel [2004], p.217, for details. Here and in what follows we use the canonical truncation function $\mathbf{1}(|x| < 1)$.

(iii) For every ω, $Q(\cdot, \omega)$ *is a measure on* I.

The integral of a function $f : I \mapsto \mathbb{R}$ *with respect to a measure* Q *is written as*

$$\int_A f(s,x)\, Q(ds, dx), \quad A \in \mathcal{B}(I).$$

Theorem 4.3 (Lévy-Itô decomposition). *Let* $X = (X_t)_{t \geq 0}$ *be a Lévy process with characteristic triplet* (γ, b^2, ν). *Then* X *can be decomposed into*

$$X_t = \gamma t + bB_t + M_t + J_t$$

with

- *a drift-term* γt *with constant* $\gamma \in \mathbb{R}$
- *a standard Brownian motion* (B_t) *scaled by a constant* $b \in \mathbb{R}$
- *a quadratic pure jump process*[6]

  $$M_t = \int_{|x|<1} x\,(Q([0,t], dx) - t\nu(dx)),$$

  where $Q(dt, dx)$ *is a time-homogeneous Poisson random measure on* $\mathbb{R}_+ \times \mathbb{R}$ *and* ν *the corresponding intensity (Lévy) measure (also called* $\mathbb{P}$*-compensator of* Q*).* (M_t) *can be regarded as a compensated (finite or countably infinite) sum of jumps with absolute size smaller than one.*
- *a compound Poisson process*

  $$J_t = \int_{|x| \geq 1} x\, Q([0,t], dx) = \sum_{s \in (0,t]}^{|\Delta X_s| \geq 1} \Delta X_s,$$

  where we use the notation $\Delta X_s = X_s - X_{s-}$ *with* $X_{s-} = \lim_{t \nearrow s} X_t$.

The processes appearing in the decomposition are independent Lévy processes, and (bB_t) *and* (M_t) *are martingales.*[7]

[6] A process (M_t) is called quadratic pure jump, if the continuous part of its quadratic variation process, $([M,M]^c_t)$, is identically zero. In this case, the quadratic variation reduces to $[M,M]_t = \sum_{s \in [0,t]}^{\Delta M_s \neq 0} (\Delta M_s)^2$.

[7] Again we note that the drift depends on the truncation function chosen in the decomposition

In case $\int_{|x|<1} |x|\, \nu(dx) < \infty$,[8] the Lévy-Itô decomposition of a Lévy process with characteristic triplet (γ, b^2, ν) simplifies to

$$X_t = \tilde{\gamma} t + b B_t + J_t$$

with

$$J_t = \int_{\mathbb{R}} x Q([0,t], dx) = \sum_{s \in (0,t]}^{\Delta X_s \neq 0} \Delta X_s$$

a compound Poisson process, and

$$\tilde{\gamma} = \gamma - \int_{|x|<1} x\, \nu(dx).$$

4.1.2 Additive processes and generalisations

When building concrete Lévy-based financial models, one soon encounters their main limitation: the stationarity of increments leads to rigid scaling properties, which are hardly ever observed in real-world financial data. This makes incorporating observed term structures (for instance of implied volatilities) an almost hopeless effort. The logical consequence of this shortcoming is to give up stationarity and allow for time-inhomogeneity, which leads to the class of *(time-)inhomogeneous Lévy processes,* also known as *additive processes.* Fortunately, additive processes are almost as tractable as Lévy processes, while providing a much greater degree of flexibility.

Definition 4.3. *Assume a filtered probability space* $(\Omega, \mathbb{F} = (\mathcal{F}_t), \mathcal{F}, \mathbb{P})$ *satisfying the usual hypotheses. A càdlàg, adapted, real-valued process* $X = (X_t)_{t \geq 0}$ *with* $X_0 = 0$ *a.s. is an* inhomogeneous Lévy process *or* additive process *if*

(i) X *has increments independent of the past, i.e.* $X_t - X_s$ *is independent of* $\mathcal{F}_s$, $0 \leq s \leq t < \infty$.

(ii) X *is continuous in probability, i.e.* $\forall \epsilon > 0$ *and* $\forall t \geq 0$ *we have* $\lim_{h \to 0} \mathbb{P}(|X_{t+h} - X_t| > \epsilon) = 0$.

[8] In the absence of a Brownian component, $\int_{|x|<1} |x|\, \nu(dx) < \infty$ holds if and only if the paths of the Lévy process are of finite variation on finite time intervals a.s.

In analogy to the time-homogeneous case, one has (see Sato [1999], p.52, or Cont & Tankov [2004], p.457)

Theorem 4.4 (Lévy-Khintchine representation for additive processes). *Let $X = (X_t)_{t\geq 0}$ be an additive process. Then X_t has an infinitely divisible distribution, and the characteristic function of X_t is*

$$\mathbb{E}\left[e^{iuX_t}\right] = e^{-\phi_t(u)}$$

with

$$\phi_t(u) = \frac{b_t^2}{2}u^2 - i\Gamma_t u + \int_{|x|\geq 1}\left(1 - e^{iux}\right)\,\mu_t(dx) + \int_{|x|<1}\left(1 - e^{iux} + iux\right)\,\mu_t(dx)$$

The so-called spot characteristics $\left(\Gamma_t, b_t^2, \mu_t\right)_{t\geq 0}$ *satisfy the following conditions:*

- μ_t *is a positive measure on* $\mathbb{R}$ *with* $\mu_t(0) = 0$ *and* $\int_{\mathbb{R}} \min\left\{x^2, 1\right\}\,\mu_t(dx) < \infty\, \forall t \geq 0$.
- $\left(b_t^2\right)$ *and* (Γ_t) *are deterministic processes.*
- $b_0 = 0,\ \mu_0 \equiv 0,\ \Gamma_0 = 0.$
- *for* $0 \leq s \leq t$, *we have* $b_t^2 \geq b_s^2$ *and* $\mu_t(B) \geq \mu_s(B)$ *for all* $B \in \mathcal{B}(\mathbb{R})$.
- *for* $s \to t$, *we have* $b_s^2 \to b_t^2,\ \Gamma_s \to \Gamma_t$ *and* $\mu_s(B) \to \mu_t(B)$ *for all* $B \in \mathcal{B}(\mathbb{R})$ *with* $B \subset \{x : |x| \geq \epsilon\}$ *for some* $\epsilon > 0$.

Conversely, for a family of triplets $\left(\Gamma_t, b_t^2, \mu_t\right)_{t\geq 0}$ *that satisfies the above conditions, there exists an additive process with* $\left(\Gamma_t, b_t^2, \mu_t\right)_{t\geq 0}$ *as spot characteristics.*

Observe that for additive processes, the exponent of the characteristic function is no longer linear in t.

The following example provides a convenient way to construct additive processes.

Example 4.1. *Consider a continuous deterministic function* $a : [0,T] \mapsto \mathbb{R}$ *with* $\int_0^T a_t^2\,dt < \infty$, *a family* $(\nu_t)_{t\in[0,T]}$ *of Lévy measures verifying* $\int_0^T \int_{\mathbb{R}} \min\left\{1, x^2\right\}\,\nu_t(dx)\,dt < \infty$, *and*

a deterministic function $\gamma : [0,T] \mapsto \mathbb{R}$ *of finite variation. Then the spot characteristics* $\left(\Gamma_t, b_t^2, \mu_t\right)_{t\in[0,T]}$ *given by*

$$\Gamma_t = \int_0^t \gamma_s \, ds$$
$$b_t^2 = \int_0^t a_s^2 \, ds$$
$$\mu_t(B) = \int_0^t \nu_s(B) \, ds \quad \forall B \in \mathcal{B}(\mathbb{R}) \quad \forall t \in [0,T]$$

define a unique additive process. $\left(\gamma_t, a_t^2, \nu_t\right)_{t\in[0,T]}$ *are called* local characteristics.

We now state a version of the Lévy-Itô decomposition for additive processes under some simplifying assumptions, compare Cont & Tankov [2004], p.452.

Theorem 4.5 (Lévy-Itô decomposition for additive processes). *Let* $X = (X_t)_{t\in[0,T]}$ *be an additive process with local characteristics* $\left(\gamma_t, a_t^2, \nu_t\right)_{t\in[0,T]}$. *Assume that the conditions*

(i) $\nu([0,t], dx) = \mu_t(dx)$ *is absolutely continuous in* t *with respect to the Lebesgue-measure, i.e. is of the form*

$$\nu([0,t], B) = \mu_t(B) = \int_0^t \nu_s(B) \, ds \quad \forall B \in \mathcal{B}(\mathbb{R}) \quad \forall t \in [0,T]$$

with a family $(\nu_t)_{t\in[0,T]}$ *of Lévy measures verifying* $\int_0^T \int_{\mathbb{R}} \min\left\{1, x^2\right\} \nu_t(dx) \, dt < \infty$

(ii) $\int_0^T \int_{|x|\geq 1} |x| \, \nu_s(dx) \, ds < \infty$

hold. Then X *can be split into*

$$X_t = \tilde{\Gamma}_t + \int_0^t a_s \, dB_s + M_t$$

with

- *a drift-term* $\tilde{\Gamma}_t = \int_0^t \gamma_s \, ds + \int_0^t \int_{|x|\geq 1} x \, \nu_s(dx) \, ds$
- *a quadratic pure jump process* (M_t) *with*

$$M_t = \int_{\mathbb{R}} x \left[Q([0,t], dx) - \nu([0,t], dx)\right] = \int_0^t \int_{\mathbb{R}} x \left[Q(ds, dx) - \nu_s(dx) \, ds\right]$$

and $Q(dt, dx)$ *a time-inhomogeneous Poisson random measure on* $[0,T] \times \mathbb{R}$ *and* ν *the corresponding Lévy measure.*

The processes appearing in the decomposition are independent Lévy processes, and $\left(\int_0^t a_s\, dB_s\right)$ *and* (M_t) *are martingales.*

The Lévy-Itô decomposition can of course be formulated under less restrictive assumptions. However, the above version is still general enough to serve our purposes.

The following definition generalises the notion of a Poisson random measure.

Definition 4.4. *A random measure on* $I = (0,\infty)\times\mathbb{R}$ *is a family* $Q = (Q(\omega; dt, dx) : \omega \in \Omega)$ *of non-negative measures on* $(I, \mathcal{B}(I))$.

In the following section, we will deal with processes more general than additive ones, which will be of the form

$$\tilde{\Gamma}_t + \int_0^t a_s\, dB_s + \int_0^t \int_{\mathbb{R}} x\,[Q(ds,dx) - \nu(ds,dx)] \tag{4.1}$$

where the compensator ν is a predictable measure with the consequence that the increments of the process are not necessarily independent any longer. That is, we will have to deal with semi-martingales with jumps of the form (4.1), which will generally have neither stationary nor independent increments. An integer-valued, non-negative random measure $Q(dt, dx)$ governs the mechanism whereby jumps occur. The compensator ν of Q is the unique predictable measure with the property that

$$Q([0,t], B) - \nu([0,t], B)$$

is a martingale for all $B \in \mathcal{B}(\mathbb{R})$. It is also possible to characterise the compensator as the unique predictable measure such that

$$\mathbb{E}\left[\int_0^t \int_B H(s,x)\,[Q(ds,dx) - \nu(ds,dx)]\right] = 0$$

for all $B \in \mathcal{B}(\mathbb{R})$ and all predictable processes H; see, for instance, Chan [1999].

We now recall Itô's formula for càdlàg semi-martingales (see e.g. Protter [2003], Chapter 2, Theorem 32, or Bingham & Kiesel [2004], Theorem 5.10.1).

Theorem 4.6 (Itô's formula). *Let* (X_t) *be a càdlàg semi-martingale, and* $f \in C^2(\mathbb{R})$.

Then $(f(X_t))$ *is again a càdlàg semi-martingale, and the following formula holds:*

$$f(X_t) = f(X_0) + \int_{0+}^{t} f'(X_{s-})\,dX_s + \frac{1}{2}\int_{0+}^{t} f''(X_{s-})\,d[X,X]_s^c + \sum_{0<s\le t}^{\Delta X_s \neq 0} \left[f'(X_s) - f'(X_{s-}) - f'(X_{s-})\Delta X_s\right].$$

The next important result essentially translates Theorems 3.24 and 5.19 of Chapter 3 of Jacod & Shiryaev [1987] to our present setting, and describes how a change of measure affects the Brownian and jump parts of a process of the form (4.1), compare also Chan [1999].

Theorem 4.7. *Let* $(G_t)_{t\in[0,T]}$ *be a predictable process,* $(H(t,x))_{t\in[0,T]}$ *be predictable for fixed* x *and the mapping* $x \mapsto H(t,x)$ *Borel-measurable for fixed* t*. Assume* $H \geq 0$ *and* $H(t,0) = 1$ *for all* $t \in [0,T]$*. Define a process* $Z = (Z_t)_{t\in[0,T]}$ *by*

$$Z_t = \exp\left\{\int_0^t G_s\,dB_s - \frac{1}{2}\int_0^t G_s^2\,ds + \int_0^t\int_{\mathbb{R}} (H(s,x)-1)(Q(ds,dx) - \nu_s(dx)\,ds)\right\} \cdot \prod_{0<s\le t}^{H(s,\Delta X_s)\neq 0} H(s,\Delta X_s)\exp\{-H(s,\Delta X_s)+1\}.$$

Then Z *is a non-negative local martingale with* $Z_0 = 1$ *and* Z *is positive if and only if* $H > 0$*. Assume* $\mathbb{E}[Z_T] = 1$.[9] *Then the measure* $\mathbb{P}^*$ *with*

$$\left.\frac{d\mathbb{P}^*}{d\mathbb{P}}\right|_{\mathcal{F}_T} = Z_T$$

is absolutely continuous with respect to $\mathbb{P}$ *on* $\mathcal{F}_T$*.*

For $(B_t)_{t\in[0,T]}$ *a* $\mathbb{P}$*-Brownian motion, the process* $(B_t^*)_{t\in[0,T]}$ *with*

$$B_t^* = B_t - \int_0^t G_s\,ds$$

is a Brownian motion under $\mathbb{P}^*$*, and for* Q *a random measure with* $\mathbb{P}$*-compensator* $\nu(dt,dx) = dt\,\nu_t(dx)$*, the* $\mathbb{P}^*$*-compensator is of the form* $\nu^*(dt,dx) = dt\,\nu_t^*(dx)$ *with*

$$\nu_t^*(dx) = H(t,x)\,\nu_t(dx).$$

Now we have all the necessary tools at hand to give a novel construction of a Lévy-based LIBOR market model.

[9] Observe that this implies that Z is a martingale.

4.2 Lévy-driven LIBOR market models

Our setup is very much the same as in Chapter 3. We are given a tenor structure $\mathcal{T} = \{T_0, \ldots, T_n\}$ as a set of maturities T_i with $0 = T_0 < T_1 < \cdots < T_n$, where T_n is the time horizon of our economy. $\mathcal{T}$ is associated with a set of $\{\tau_1, \ldots, \tau_n\}$ of year fractions, where $\tau_i = T_i - T_{i-1}$, $i = 1, \ldots, n$. For simplicity, we assume $\tau_i \equiv \delta$. Moreover, we assume that in the financial market under consideration, there exist zero-coupon bonds $p(\cdot, T_i)$ of all maturities T_i, $i = 1, \ldots, n$. The discretely compounded forward LIBOR rate prevailing at time t over the future period from T_{i-1} to T_i is

$$L(t, T_{i-1}) = \frac{p(t, T_{i-1}) - p(t, T_i)}{\delta p(t, T_i)}, \quad 0 \leq t \leq T_{i-1}.$$

Our model is built on a stochastic basis $\left(\Omega, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T_{n-1}]}, \mathcal{F}, \mathbb{P}^{T_n}\right)$ satisfying the usual conditions, on which an additive process $L^{T_n} = \left(L_t^{T_n}\right)_{t \in [0, T_{n-1}]}$ is defined. We assume that $\mathbb{P}^{T_n}$ is the T_n-forward measure. As we already know, the forward LIBOR process $(L(t, T_{n-1}))$ must be a $\mathbb{P}^{T_n}$-local martingale. Now we take an approach that deviates from the one we followed in the diffusion setting: We do *not* posit an SDE for the LIBOR dynamics. Rather, we describe the dynamics directly by an *exponential additive model* by postulating

$$\begin{aligned} L(t, T_{n-1}) =& L(0, T_{n-1}) \exp\left(\int_0^t \lambda(s, T_{n-1})\, dL_s^{T_n}\right) \\ =& L(0, T_{n-1}) \exp\left(X_t^{T_{n-1}}\right), \quad t \in [0, T_{n-1}] \end{aligned}$$

with $L(0, T_{n-1}) > 0$, and a deterministic function $\lambda(\cdot, T_{n-1}) : [0, T_{n-1}] \mapsto \mathbb{R}_+$ which is bounded by a constant $M^{T_{n-1}} \in \mathbb{R}_+$. $\lambda(\cdot, T_{n-1})$ describes the term structure of volatility of $L(\cdot, T_{n-1})$. The definition of $X_t^{T_{n-1}}$ is obvious.[10]

[10] A naïve replication of the SDE-approach taken in Chapter 3 would – in the presence of jumps smaller than -1 – lead to negative LIBOR rates. Special care would have to be taken to avoid this, for instance by introducing stochastic λ's which would extremely complicate the modelling process, or by severely restricting the class of driving processes by excluding those with jumps smaller than -1. Both approaches are unsatisfactory.

We posit[11] that the driving additive process L^{T_n} is parameterised as in Theorem 4.5:[12]

$$L_t^{T_n} = \int_0^t \gamma_s^{T_n}\, ds + \int_0^t a_s^{T_n}\, dB_s^{T_n} + \int_0^t \int_{\mathbb{R}} x \left[Q(ds,dx) - \nu_s^{T_n}(dx)\, ds\right],$$

assuming that $\left(a_s^{T_n}\right)$ and $\left(\nu_s^{T_n}\right)$ are deterministic, and that all integrals, here and in what follows, are well defined. As we will soon see, $\left(\gamma_s^{T_n}\right)$ is also deterministic.

$X^{T_{n-1}}$ is an additive process (which motivates the name exponential additive model), since

$$\begin{aligned}
X_t^{T_{n-1}} &= \int_0^t \lambda(s,T_{n-1})\, dL_s^{T_n} \\
&= \int_0^t \lambda(s,T_{n-1})\gamma_s^{T_n}\, ds + \int_0^t \lambda(s,T_{n-1})a_s^{T_n}\, dB_s^{T_n} \\
&\quad + \int_0^t \int_{\mathbb{R}} \lambda(s,T_{n-1})x \left[Q(ds,dx) - \nu_s^{T_n}(dx)\, ds\right] \\
&= \int_0^t \lambda(s,T_{n-1})\gamma_s^{T_n}\, ds + \int_0^t \lambda(s,T_{n-1})a_s^{T_n}\, dB_s^{T_n} \\
&\quad + \int_0^t \int_{\mathbb{R}} x \left[Q\left(ds, \frac{dx}{\lambda(s,T_{n-1})}\right) - \nu_s^{T_n}\left(ds, \frac{dx}{\lambda(s,T_{n-1})}\right)\right].
\end{aligned}$$

Now we apply Itô's formula to derive the dynamics of $(L(t,T_{n-1}))$.

$$\begin{aligned}
&\exp\left(X_t^{T_{n-1}}\right) \\
&= 1 + \int_{0+}^t \exp\left(X_{s-}^{T_{n-1}}\right) dX_s^{T_{n-1}} + \frac{1}{2}\int_{0+}^t \exp\left(X_{s-}^{T_{n-1}}\right) d[X^{T_{n-1}}, X^{T_{n-1}}]_s^c \\
&\quad + \sum_{0<s\le t}^{\Delta X_s^{T_{n-1}}\neq 0} \left[\exp\left(X_s^{T_{n-1}}\right) - \exp\left(X_{s-}^{T_{n-1}}\right) - \exp\left(X_{s-}^{T_{n-1}}\right)\Delta X_s^{T_{n-1}}\right] \\
&= 1 + \int_{0+}^t \exp\left(X_{s-}^{T_{n-1}}\right) dX_s^{T_{n-1}} + \frac{1}{2}\int_{0+}^t \exp\left(X_{s-}^{T_{n-1}}\right) \lambda(s,T_{n-1})^2 \left(a_s^{T_n}\right)^2 ds \\
&\quad + \sum_{0<s\le t}^{\Delta X_s^{T_{n-1}}\neq 0} \left[\exp\left(X_{s-}^{T_{n-1}} + \Delta X_s^{T_{n-1}}\right) - \exp\left(X_{s-}^{T_{n-1}}\right) - \exp\left(X_{s-}^{T_{n-1}}\right)\Delta X_s^{T_{n-1}}\right] \\
&= 1 + \int_{0+}^t \exp\left(X_{s-}^{T_{n-1}}\right) dX_s^{T_{n-1}} + \frac{1}{2}\int_{0+}^t \exp\left(X_{s-}^{T_{n-1}}\right) \lambda(s,T_{n-1})^2 \left(a_s^{T_n}\right)^2 ds \\
&\quad + \int_{0+}^t \int_{\mathbb{R}} \exp\left(X_{s-}^{T_{n-1}}\right)\left(e^x - 1 - x\right) Q\left(ds, \frac{dx}{\lambda(s,T_{n-1})}\right)
\end{aligned}$$

[11] Compare also Björk *et al.* [1997], p.151, who use a similar approach.

[12] $\int_0^t \gamma_s^{T_n}\, ds$ here corresponds to $\tilde{\Gamma}_t$ in Theorem 4.5.

$$= 1 + \int_{0+}^{t} \exp\left(X_{s-}^{T_{n-1}}\right) dX_s^{T_{n-1}} + \frac{1}{2}\int_{0+}^{t} \exp\left(X_{s-}^{T_{n-1}}\right) \lambda(s,T_{n-1})^2 \left(a_s^{T_n}\right)^2 ds$$
$$+ \int_{0+}^{t}\int_{\mathbb{R}} \exp\left(X_{s-}^{T_{n-1}}\right)\left(e^{\lambda(s,T_{n-1})x} - 1 - \lambda(s,T_{n-1})x\right) Q(ds,dx).$$

In differential form, we get

$$\frac{d\exp\left(X_t^{T_{n-1}}\right)}{\exp\left(X_{t-}^{T_{n-1}}\right)}$$
$$= dX_t^{T_{n-1}} + \frac{1}{2}\lambda(t,T_{n-1})^2\left(a_t^{T_n}\right)^2 dt + \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - 1 - \lambda(t,T_{n-1})x\right) Q(dt,dx)$$
$$= \lambda(t,T_{n-1})\gamma_t^{T_n}\,dt + \lambda(t,T_{n-1})a_t^{T_n}\,dB_t^{T_n} + \int_{\mathbb{R}} \lambda(t,T_{n-1})x\left[Q(dt,dx) - \nu_t^{T_n}(dx)\,dt\right]$$
$$+ \frac{1}{2}\lambda(t,T_{n-1})^2\left(a_t^{T_n}\right)^2 dt + \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - 1 - \lambda(t,T_{n-1})x\right) Q(dt,dx)$$
$$= \lambda(t,T_{n-1})\gamma_t^{T_n}\,dt + \lambda(t,T_{n-1})a_t^{T_n}\,dB_t^{T_n} + \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - \lambda(t,T_{n-1})x - 1\right)\nu_t^{T_n}(dx)\,dt$$
$$+ \frac{1}{2}\lambda(t,T_{n-1})^2\left(a_t^{T_n}\right)^2 dt + \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - 1\right)\left[Q(dt,dx) - \nu_t^{T_n}(dx)\,dt\right].$$

We have the semi-martingale decomposition[13]

$$\exp\left(X_t^{T_{n-1}}\right) = M_t + V_t,$$

where (M_t) is a martingale with

$$dM_t = \exp\left(X_{t-}^{T_{n-1}}\right)\left[\lambda(t,T_{n-1})a_t^{T_n}\,dB_t^{T_n} + \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - 1\right)\left[Q(dt,dx) - \nu_t^{T_n}(dx)\,dt\right]\right]$$

and (V_t) is a process of finite variation with

$$dV_t = \exp\left(X_{t-}^{T_{n-1}}\right)\left[\lambda(t,T_{n-1})\gamma_t^{T_n}\,dt + \frac{1}{2}\lambda(t,T_{n-1})^2\left(a_t^{T_n}\right)^2 dt\right.$$
$$\left. + \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - \lambda(t,T_{n-1})x - 1\right)\nu_t^{T_n}(dx)\,dt\right].$$

Now we have to make sure that $\left(\exp\left(X_t^{T_{n-1}}\right)\right)$ (and thus $(L(t,T_{n-1}))$) is a martingale under $\mathbb{P}^{T_n}$. We achieve this by demanding that $dV_t = 0$ $\mathbb{P}^{T_n}$-almost surely for all $t \geq 0$, which leads to

$$\lambda(t,T_{n-1})\gamma_t^{T_n} + \frac{1}{2}\lambda(t,T_{n-1})^2\left(a_t^{T_n}\right)^2 + \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - \lambda(t,T_{n-1})x - 1\right)\nu_t^{T_n}(dx) = 0.$$

[13] Compare e.g. Cont & Tankov [2004], p.284.

Solving for $\gamma_t^{T_n}$ then gives

$$\gamma_t^{T_n} = -\frac{1}{2}\lambda(t,T_{n-1})\left(a_t^{T_n}\right)^2 - \frac{1}{\lambda(t,T_{n-1})}\int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - \lambda(t,T_{n-1})x - 1\right)\nu_t^{T_n}(dx),$$

and all parameters are uniquely determined. Using this deterministic (thus predictable) process $\left(\gamma_t^{T_n}\right)$, we find $V_t = V_0$ for all $t \geq 0$, and thus

$$\exp\left(X_t^{T_{n-1}}\right) = M_t + V_0,$$

or

$$\begin{aligned} d\exp\left(X_t^{T_{n-1}}\right) =& \exp\left(X_{t-}^{T_{n-1}}\right)\Big[\lambda(t,T_{n-1})a_t^{T_n}\,dB_t^{T_n} \\ &+ \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - 1\right)\Big[Q(dt,dx) - \nu_t^{T_n}(dx)\,dt\Big]\Big]. \end{aligned}$$

This translates to the following exponential SDE for the LIBOR-dynamics:

$$\begin{aligned} dL(t,T_{n-1}) =& L(t-,T_{n-1})\Big[\lambda(t,T_{n-1})a_t^{T_n}\,dB_t^{T_n} \\ &+ \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x} - 1\right)\Big[Q(dt,dx) - \nu_t^{T_n}(dx)\,dt\Big]\Big], \end{aligned}$$

and apparently, $(L(t,T_{n-1}))$ is a stochastic (or Doléans-Dade) exponential and as such a martingale under $\mathbb{P}^{T_n}$.

As we know from Chapter 3, the T_{n-1}-forward measure $\mathbb{P}^{T_{n-1}}$ is the martingale measure with respect to the numéraire $p(\cdot,T_{n-1})$, and, as already shown, the LIBOR process $(L(t,T_{n-2}))$ has to be a $\mathbb{P}^{T_{n-1}}$ martingale. Thus, an application of the change of numéraire theorem (see Geman *et al.* [1995] or Bingham & Kiesel [2004], p.239) leads to

$$Z_t^{T_{n-1}} := \frac{d\mathbb{P}^{T_{n-1}}}{d\mathbb{P}^{T_n}}\bigg|_{\mathcal{F}_t} = \frac{1+\delta L(t,T_{n-1})}{1+\delta L(0,T_{n-1})}, \quad t \in [0,T_{n-2}].$$

The dynamics of the Radon-Nikodým density process $\left(Z_t^{T_{n-1}}\right)$ are

$$\begin{aligned} dZ_t^{T_{n-1}} =& d\left(\frac{1+\delta L(t,T_{n-1})}{1+\delta L(0,T_{n-1})}\right) \\ =& \frac{\delta}{1+\delta L(0,T_{n-1})}\,dL(t,T_{n-1}) \\ =& \frac{\delta L(t-,T_{n-1})}{1+\delta L(0,T_{n-1})}\Big[\lambda(t,T_{n-1})a_t^{T_n}\,dB_t^{T_n} \end{aligned}$$

$$
\begin{aligned}
&+\int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x}-1\right)\left[Q(dt,dx)-\nu_t^{T_n}(dx)\,dt\right]\Big] \\
&=\frac{1+\delta L(t-,T_{n-1})}{1+\delta L(0,T_{n-1})}\frac{\delta L(t-,T_{n-1})}{1+\delta L(t-,T_{n-1})}\Big[\lambda(t,T_{n-1})a_t^{T_n}\,dB_t^{T_n} \\
&\quad+\int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1})x}-1\right)\left[Q(dt,dx)-\nu_t^{T_n}(dx)\,dt\right]\Big] \\
&=Z_{t-}^{T_{n-1}}\left[G_t^{T_{n-1}}\,dB_t^{T_n}+\int_{\mathbb{R}}\left(H^{T_{n-1}}(t,x)-1\right)\left[Q(dt,dx)-\nu_t^{T_n}(dx)\,dt\right]\right], \qquad (4.2)
\end{aligned}
$$

with

$$
G_t^{T_{n-1}}=\frac{\delta L(t-,T_{n-1})}{1+\delta L(t-,T_{n-1})}\lambda(t,T_{n-1})a_t^{T_n}
$$

a predictable process, and

$$
H^{T_{n-1}}(t,x)-1=\frac{\delta L(t-,T_{n-1})}{1+\delta L(t-,T_{n-1})}\left(e^{\lambda(t,T_{n-1})x}-1\right)
$$

a process with the property that it is predictable for fixed x and the function $x\mapsto H^{T_{n-1}}(t,x)$ is Borel-measurable for fixed t. Furthermore, we recognise $\left(Z_t^{T_{n-1}}\right)$ as a Doléans-Dade exponential with[14]

$$
\begin{aligned}
Z_t^{T_{n-1}}=&\exp\left\{\int_0^t G_s^{T_{n-1}}\,dB_s^{T_n}-\frac{1}{2}\int_0^t\left(G_s^{T_{n-1}}\right)^2\,ds\right. \\
&\left.+\int_0^t\int_{\mathbb{R}}\left(H^{T_{n-1}}(s,x)-1\right)\left[Q(ds,dx)-\nu_s^{T_n}(dx)\,ds\right]\right\} \\
&\cdot\prod_{0<s\le t}^{H^{T_{n-1}}(s,\Delta X_s)\neq 0}H^{T_{n-1}}(s,\Delta X_s)\exp\left\{-H^{T_{n-1}}(s,\Delta X_s)+1\right\},
\end{aligned}
$$

where by convention the empty product is equal to 1. So $Z_0^{T_{n-1}}=1$, and a closer look at the definition of H reveals that H is strictly positive, and thus $Z^{T_{n-1}}$ is a strictly positive martingale, as we would have expected from its definiton as Radon-Nikodým density process. We are now exactly in the situation of Theorem 4.7, which tells us that

$$
B_t^{T_{n-1}}=B_t^{T_n}-\int_0^t G_s^{T_{n-1}}\,ds
$$

is a Brownian motion under $\mathbb{P}^{T_{n-1}}$, and

$$
\nu_t^{T_{n-1}}(dx)=H^{T_{n-1}}(t,x)\,\nu_t^{T_n}(dx)
$$

[14]For general solutions to SDEs of the form (4.2), see Cont & Tankov [2004] or Protter [2003].

is the $\mathbb{P}^{T_{n-1}}$-compensator of $Q(dt, dx)$.

Now we posit that the dynamics of $(L(t, T_{n-2}))$ under $\mathbb{P}^{T_{n-1}}$ are given by

$$\begin{aligned} L(t, T_{n-2}) =& L(0, T_{n-2}) \exp\left(\int_0^t \lambda(s, T_{n-2})\, dL_s^{T_{n-1}} \right) \\ =& L(0, T_{n-2}) \exp\left(X_t^{T_{n-2}} \right), \quad t \in [0, T_{n-2}], \end{aligned}$$

with $L(0, T_{n-2}) > 0$. In analogy to above, we assume that

$$L_t^{T_{n-1}} = \int_0^t \gamma_s^{T_{n-1}}\, ds + \int_0^t a_s^{T_{n-1}}\, dB_s^{T_{n-1}} + \int_0^t \int_{\mathbb{R}} x \left[Q(ds, dx) - \nu_s^{T_{n-1}}(dx)\, ds \right], \qquad (4.3)$$

with deterministic $\left(a_s^{T_{n-1}}\right)$; that is, we assume that $(L(t, T_{n-2}))$ is driven by the same noise as $(L(t, T_{n-1}))$.

At this point, we need to stress that (4.3) is no additive process anymore, since the compensators $\left(\nu_t^{T_{n-1}}\right)$ are obviously stochastic, as they depend on the realisation of the LIBORs. Therefore, $(L(t, T_{n-2}))$ will in general lack the independent increments property.

Proceeding as above, we have

$$\begin{aligned} X_t^{T_{n-2}} &= \int_0^t \lambda(s, T_{n-2})\, dL_s^{T_{n-1}} \\ &= \int_0^t \lambda(s, T_{n-2}) \gamma_s^{T_{n-1}}\, ds + \int_0^t \lambda(s, T_{n-2}) a_s^{T_{n-1}}\, dB_s^{T_{n-1}} \\ &\quad + \int_0^t \int_{\mathbb{R}} x \left[Q\left(ds, \frac{dx}{\lambda(s, T_{n-2})} \right) - \nu_s^{T_{n-1}}\left(\frac{dx}{\lambda(s, T_{n-2})} \right) ds \right]. \end{aligned}$$

Imposing the condition that $(L(t, T_{n-2}))$ be a $\mathbb{P}^{T_{n-1}}$-martingale and following the same arguments as above implies the condition

$$\begin{aligned} \gamma_t^{T_{n-1}} =& -\frac{1}{2} \lambda(t, T_{n-2}) \left(a_t^{T_{n-1}} \right)^2 \\ & - \frac{1}{\lambda(t, T_{n-2})} \int_{\mathbb{R}} \left(e^{\lambda(t, T_{n-2}) x} - \lambda(s, T_{n-2}) x - 1 \right) \nu_t^{T_{n-1}}(dx) \quad \mathbb{P}^{T_{n-1}}\text{-a.s.} \end{aligned} \qquad (4.4)$$

for the drift, which is, due to the stochastic nature of the compensator, also stochastic, and in addition predictable, since $\left(\nu_t^{T_{n-1}}\right)$ is predictable. Again, with (4.4), all parameters in (4.3) are uniquely determined. With this specification of $\left(\gamma_t^{T_{n-1}}\right)$, the martingale

dynamics of $(L(t,T_{n-2}))$ under $\mathbb{P}^{T_{n-1}}$ are

$$\begin{aligned} dL(t,T_{n-2}) =& L(t-,T_{n-2})\left[\lambda(t,T_{n-2})a_t^{T_{n-1}}\,dB_t^{T_{n-1}}\right. \\ &\left. + \int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-2})x}-1\right)\left[Q(dt,dx)-\nu_t^{T_{n-1}}(dx)\,dt\right]\right]. \end{aligned}$$

We find it instructive to carry out one more step of our inductive procedure. By the same reasoning as above, an application of the change of numéraire theorem provides the Radon-Nikodým density of the T_{n-2}-forward measure $\mathbb{P}^{T_{n-2}}$ with respect to $\mathbb{P}^{T_{n-1}}$ as

$$Z_t^{T_{n-2}} := \left.\frac{d\mathbb{P}^{T_{n-2}}}{d\mathbb{P}^{T_{n-1}}}\right|_{\mathcal{F}_t} = \frac{1+\delta L(t,T_{n-2})}{1+\delta L(0,T_{n-2})}, \quad t\in[0,T_{n-3}].$$

The dynamics of the Radon-Nikodým density process $\left(Z_t^{T_{n-2}}\right)$ are

$$dZ_t^{T_{n-2}} = Z_{t-}^{T_{n-2}}\left[G_t^{T_{n-2}}\,dB_t^{T_{n-1}} + \int_{\mathbb{R}}\left(H^{T_{n-2}}(t,x)-1\right)\left[Q(dt,dx)-\nu_t^{T_{n-1}}(dx)\,dt\right]\right],$$

with

$$G_t^{T_{n-2}} = \frac{\delta L(t-,T_{n-2})}{1+\delta L(t-,T_{n-2})}\lambda(t,T_{n-2})a_t^{T_{n-1}}$$

a predictable process, and

$$H^{T_{n-2}}(t,x)-1 = \frac{\delta L(t-,T_{n-2})}{1+\delta L(t-,T_{n-2})}\left(e^{\lambda(t,T_{n-2})x}-1\right)$$

a process with the property that it is predictable for fixed x and the function $x\mapsto H^{T_{n-2}}(t,x)$ is Borel-measurable for fixed t. $\left(Z_t^{T_{n-2}}\right)$ is a Doléans-Dade exponential with

$$\begin{aligned} Z_t^{T_{n-2}} =& \exp\left\{\int_0^t G_s^{T_{n-2}}\,dB_s^{T_{n-1}} - \frac{1}{2}\int_0^t\left(G_s^{T_{n-2}}\right)^2 ds\right. \\ &\left. + \int_0^t\int_{\mathbb{R}}\left(H^{T_{n-2}}(s,x)-1\right)\left[Q(ds,dx)-\nu_s^{T_{n-1}}(dx)\,ds\right]\right\} \\ &\cdot \prod_{0<s\le t}^{H^{T_{n-2}}(s,\Delta X_s)\neq 0} H^{T_{n-2}}(s,\Delta X_s)\exp\left\{-H^{T_{n-2}}(s,\Delta X_s)+1\right\}. \end{aligned}$$

We are now in the position to apply Theorem 4.7, and conclude that

$$\begin{aligned} B_t^{T_{n-2}} =& B_t^{T_{n-1}} - \int_0^t G_s^{T_{n-2}}\,ds \\ =& B_t^{T_n} - \int_0^t G_s^{T_{n-1}}\,ds - \int_0^t G_s^{T_{n-2}}\,ds \end{aligned}$$

is a Brownian motion under $\mathbb{P}^{T_{n-2}}$, and

$$\begin{aligned}\nu_t^{T_{n-2}}(dx) &= H^{T_{n-2}}(t,x)\nu_t^{T_{n-1}}(dx)\\ &= H^{T_{n-2}}(t,x)H^{T_{n-1}}(t,x)\nu_t^{T_n}(dx)\end{aligned}$$

is the $\mathbb{P}^{T_{n-2}}$-compensator of $Q(dt,dx)$.

The following central theorem summarises our findings.

Theorem 4.8 (LIBOR dynamics in a Lévy-based LIBOR market model). *Let the above assumptions hold true. Then the dynamics of the LIBOR process* $(L(t,T_{n-1-k}))$, $k \in \{0,\dots,n-2\}$, *are described by the SDE*

$$\begin{aligned}dL(t,T_{n-1-k}) =& L(t-,T_{n-1-k})\Big[\lambda(t,T_{n-1-k})a_t^{T_{n-k}}\,dB_t^{T_{n-k}}\\ &+\int_{\mathbb{R}}\left(e^{\lambda(t,T_{n-1-k})x}-1\right)\Big[Q(dt,dx)-\nu_t^{T_{n-k}}(dx)\,dt\Big]\Big]\end{aligned} \tag{4.5}$$

with

$$B_t^{T_{n-k}} = B_t^{T_n} - \sum_{i=1}^{k}\int_0^t G_s^{T_{n-i}}\,ds$$

a $\mathbb{P}^{T_{n-k}}$*-Brownian motion, and*

$$\nu_t^{T_{n-k}}(dx) = \nu_t^{T_n}(dx)\prod_{i=1}^{k}H^{T_{n-i}}(t,x)$$

the $\mathbb{P}^{T_{n-k}}$*-compensator of* $Q(dt,dx)$.

As the proof is just an inductive application of the arguments used above, it is omitted.[15]

The above theorem is vital for the implementation of our model, as it allows us to simulate all LIBOR-rates under *one* measure, which in this case is the terminal measure $\mathbb{P}^{T_n}$.

4.3 Pricing caplets and floorlets in a Lévy LMM

The pricing of caplets and floorlets in Lévy-driven LMMs does not always entail time-consuming simulation procedures. In this section, we outline two methods that reduce

[15]The above result can also be derived by following the arguments of Eberlein & Özkan [2004].

the caplet-pricing problem to integral evaluations. It goes without saying that the same methods also apply to floorlets. This section adapts the corresponding section in Schoutens [2003] to the interest rate context.

A caplet with maturity T_k and strike K pays the holder $\delta(L(T_k,T_k)-K)^+$ at time T_{k+1}. Its value at time $t=0$ is therefore

$$C(0,T_k,K)=\delta p(0,T_{k+1})\mathbb{E}_{\mathbb{P}^{T_{k+1}}}\left[(L(T_k,T_k)-K)^+\right].$$

4.3.1 Pricing by means of the density function

If the density function of $L(T_k,T_k)$ under $\mathbb{P}^{T_{k+1}}$, say $f(x)$, is known, we can work out the caplet price by (possibly numerical) integration:

$$C(0,T_k,K)=\delta p(0,T_{k+1})\int_K^\infty (x-K)f(x)\,dx.$$

4.3.2 Pricing by means of the characteristic function

If the characteristic function of the logarithm of the LIBOR rate,

$$\phi(u)=\mathbb{E}_{\mathbb{P}^{T_{k+1}}}\left[\exp\left(iu\log L(T_k,T_k)\right)\right],$$

is known, then, according to Bakshi & Madan [2000], we can represent the caplet-price as

$$C(0,T_k,K)=\delta p(0,T_{k+1})\left(L(0,T_k)\Pi_1-K\Pi_2\right)$$

with

$$\Pi_1=\frac{1}{2}+\frac{1}{\pi}\int_0^\infty \mathrm{Re}\left(\frac{\exp(-iu\log K)\phi(u-i)}{iu}\right)du$$

$$\Pi_2=\frac{1}{2}+\frac{1}{\pi}\int_0^\infty \mathrm{Re}\left(\frac{\exp(-iu\log K)\phi(u)}{iu}\right)du.$$

Similarly to the Black-case, Π_1 is the delta of the option, while Π_2 is the probability of finishing in the money.

4.4 Approximations

4.4.1 Approximate LIBOR dynamics

As Theorem 4.8 shows, the dynamics of the LIBOR rates are quite involved. The principal problems when it comes to implementation are

- stochastic drifts of the various Brownian motions,
- stochastic compensators,
- stochastic drift terms (γ_i).

These points render the implementation not only complicated, but also inefficient, since in the simulation procedure, only relatively small step-sizes can be chosen, as one has to account for the stochastic and time-dependent nature of the above quantities. Quite naturally, one looks for simplifications that do not distort the spirit of the model too much, while facilitating and speeding up its implementation and reducing its computational burden. Inspired by an approximation technique – commonly dubbed 'freezing the drift' – that has been successfully used in the standard LMM-context for quite a few years, we apply a similar idea to our present setting.[16]

Theorem 4.9 (Approximate LIBOR dynamics under the terminal measure). *The dynamics of the LIBOR process* $(L(t,T_{n-1-k}))$, $k \in \{0,\dots,n-2\}$, *under the terminal measure* $\mathbb{P}^{T_n}$ *can be approximated by*

$$\begin{aligned} dL(t,T_{n-1-k}) =& L(t-,T_{n-1-k}) \Big[\lambda(t,T_{n-1-k}) a_t^{T_{n-k}}\, d\bar{B}_t^{T_{n-k}} \\ &+ \int_{\mathbb{R}} \left(e^{\lambda(t,T_{n-1-k})x} - 1\right) \Big[Q(dt,dx) - \bar{\nu}_t^{T_{n-k}}(dx)\, dt\Big]\Big] \end{aligned} \tag{4.6}$$

with

$$\bar{B}_t^{T_{n-k}} = B_t^{T_n} - \sum_{i=1}^{k} \int_0^t \bar{G}_s^{T_{n-i}}\, ds,$$

[16] See also Chapter 3, where we applied a related technique to derive an approximation formula for swaption prices.

where $\left(B_t^{T_n}\right)$ *is a* $\mathbb{P}^{T_n}$ *-Brownian motion, and with*

$$\bar{\nu}_t^{T_{n-k}}(dx) = \nu_t^{T_n}(dx) \prod_{i=1}^{k} \bar{H}^{T_{n-i}}(t,x), \tag{4.7}$$

where $\nu_t^{T_n}$ *is the* $\mathbb{P}^{T_n}$ *-compensator of* Q*. Furthermore,*

$$\bar{G}_s^{T_{n-i}} = \frac{\delta L(0,T_{n-i})}{1+\delta L(0,T_{n-i})} \lambda(s,T_{n-i}) a_s^{T_{n-i+1}}$$

and

$$\bar{H}^{T_{n-i}}(t,x) - 1 = \frac{\delta L(0,T_{n-i})}{1+\delta L(0,T_{n-i})} \left(e^{\lambda(t,T_{n-i})x} - 1\right) \tag{4.8}$$

are deterministic processes.

The above simplifications lead to deterministic compensators, thus independent increments, and therefore to all LIBOR processes being additive Lévy processes. The drift processes (which appear in the integrated form of the LIBOR dynamics) become

$$\begin{aligned} \bar{\gamma}_t^{T_{n-k}} = &-\frac{1}{2}\lambda(t,T_{n-1-k}) \left(a_t^{T_{n-k}}\right)^2 \\ &- \frac{1}{\lambda(t,T_{n-1-k})} \int_{\mathbb{R}} \left(e^{\lambda(t,T_{n-1-k})x} - \lambda(t,T_{n-1-k})x - 1\right) \bar{\nu}_t^{T_{n-k}}(dx). \end{aligned} \tag{4.9}$$

Summing up, we have got rid of all the problems mentioned above. But, as always in life, there's no such thing as a free lunch, and also our simplifying assumptions have their price, which comes in the form of a violation of the no-arbitrage condition. The LIBOR processes defined above (apart from $(L(t,T_{n-1}))$ will in general cease to be martingales under their respective forward measures. Even though we also adapt the drifts (γ_t) to our new situation, the martingale-condition will be violated. The reason becomes clear if we recall the construction of the drifts, where we used that the compensated jump-parts are martingales under the respective forward measures. Making the compensator deterministic (as we do above) implies that the resulting compensated jump-parts will no longer be martingales. As we do not correct the drifts for this (because it would make them stochastic), we wind up with LIBOR processes that are devoid of the martingale property under their respective forward measures.

One might legitimately ask whether an approximation that introduces arbitrage opportunities makes sense at all. In view of this, let us give two justifications of our approach. Firstly, the simulation procedures proposed in the literature for standard LMMs are not arbitrage-free either, because the very discretization of the underlying continuous time processes results in arbitrage opportunities. However, the violation of the no-arbitrage-condition due to discretization is deemed negligible (see Brigo & Mercurio [2001], p.236). Secondly, the main purpose of an LMM is to price interest rate derivatives. As simulation experiments for standard LMMs show, 'freezing the drift' has only a minor impact on derivatives prices for (almost) all sensible parameterisations,[17] and we can expect this property to carry over to our current setting. Of course, simulation studies have to be conducted to corroborate this conjecture. This is the purpose of the following section.

4.5 Implementation: A worked example

4.5.1 Implementation of the approximate model

In this part, we demonstrate step by step how to implement the approximate LMM based on formulae (4.6) through (4.8) in a concrete setup. In order to keep the exposition as simple as possible, we assume a pure jump process as driving noise. This assumption does not constitute a major limitation, since, as argued in Geman *et al.* [2001] (see also Schoutens [2003], p. 76), a realistic model for the price process of a financial asset requires a jump component, while a diffusion component is dispensable. Pure jump Lévy models can capture both (relatively rare) large jumps and (relatively frequent) smaller moves in price processes. The empirical performance of pure jump models normally cannot be enhanced by adding a diffusion component.

As for the concrete parameters, we choose $\delta = 1$, $n = 10$, a flat initial term structure $L(0,T_i) = 0.05$, $T_i \in \{0,\ldots,9\}$, and $\lambda(\cdot,\cdot) \equiv 1$.

As driving Lévy process, we take a symmetric Variance Gamma process with drift (see e.g. Madan *et al.* [1998], Schoutens [2003] or Cont & Tankov [2004] for details on Variance

[17] But see also Joshi & Rebonato [2001] for possible pitfalls.

Gamma processes), which takes the form

$$L_t^{T_{10}} = \int_0^t \bar{\gamma}_s^{T_{10}}\, ds + \tilde{L}_t^{T_{10}},$$

where

$$\tilde{L}_t^{T_{10}} = \int_0^t \int_{\mathbb{R}} x \left[Q(ds, dx) - \nu_s^{T_{10}}(dx)\, ds\right]$$

is a symmetric, driftless Variance Gamma process under $\mathbb{P}^{T_{10}}$ and as such a $\mathbb{P}^{T_{10}}$-martingale.

We assume the symmetric Lévy measure to be given by

$$\nu^{T_{10}}(dx) = \nu_s^{T_{10}}(dx) = \begin{cases} C \exp\{-G|x|\}\, |x|^{-1}\, dx & x \neq 0 \\ 0 & x = 0 \end{cases}, \quad s \geq 0,$$

with $C > 0$ and $G > 0$; in particular, $\left(\tilde{L}_t^{T_{10}}\right)$ is a time-homogeneous Lévy process. $\tilde{L}_1^{T_{10}}$ is $VG(C, G)$ distributed, and its characteristic function reads

$$\phi_{VG}(u; C, G) = \left(\frac{G^2}{G^2 + u^2}\right)^C.$$

Furthermore, using the time-homogeneity of $\left(\tilde{L}_t^{T_{10}}\right)$, we find $\tilde{L}_{t+s}^{T_{10}} - \tilde{L}_t^{T_{10}} \sim VG(sC, G)$. By virtue of the symmetry of the given Lévy measure, we find

$$\int_{\mathbb{R}} x \nu^{T_{10}}(dx) = 0$$

and thus

$$\tilde{L}_t^{T_{10}} = \int_0^t \int_{\mathbb{R}} x\, Q(dx, ds).$$

We choose $C = 2$ and $G = 12$. The drifts can now be calculated using formula (4.9), which can be evaluated to give

$$\bar{\gamma}_t^{T_{10}} = \bar{\gamma}^{T_{10}} = -\int_{\mathbb{R}} (e^x - x - 1)\, \nu^{T_{10}}(dx) = -0.013937.$$

The (exact) LIBOR-dynamics are described by

$$L(t, T_9) = L(0, T_9) \exp\left\{\tilde{L}_t^{T_{10}} - 0.013937t\right\},$$

which is a $\mathbb{P}^{T_{10}}$-martingale.

Now we will pin down the dynamics of $L(t,T_8)$. First, we observe that by formula (4.8)

$$\bar{H}^{T_9}(t,x) - 1 = \bar{H}^{T_9}(x) - 1 = \frac{0.05}{1.05}\left(e^x - 1\right),$$

and by formula (4.7)

$$\bar{\nu}^{T_9}(dx) = \nu^{T_{10}}(dx)\bar{H}^{T_9}(x),$$

and thus

$$\bar{\gamma}_t^{T_9} = \bar{\gamma}^{T_9} = -\int_{\mathbb{R}} \left(e^x - x - 1\right)\bar{\nu}^{T_9}(dx) = -0.013961.$$

Continuing,

$$\begin{aligned}
\tilde{L}_t^{T_9} =& \int_0^t \int_{\mathbb{R}} x\left[Q(ds,dx) - \bar{\nu}^{T_9}(dx)\,ds\right] \\
=& \tilde{L}_t^{T_{10}} - t\int_{\mathbb{R}} x\bar{\nu}^{T_9}(dx) \\
=& \tilde{L}_t^{T_{10}} - tk^{T_9} \\
=& \tilde{L}_t^{T_{10}} - 0.001332t
\end{aligned}$$

with an evident definition of k^{T_9}. The deterministic process $\left(tk^{T_9}\right)$ approximately accounts for the difference in the drift of the driving Lévy process under the measures $\mathbb{P}^{T_{10}}$ and $\mathbb{P}^{T_9}$; recall that $\left(\tilde{L}_t^{T_9}\right)$ is an approximation for the $\mathbb{P}^{T_9}$-martingale $\left(L_t^{T_9}\right)$ with

$$L_t^{T_9} = \int_0^t \int_{\mathbb{R}} x\left[Q(ds,dx) - \nu^{T_9}(dx)\,ds\right].$$

Summing up, we get

$$L(t,T_8) = L(0,T_8)\exp\left\{\tilde{L}_t^{T_{10}} + t\bar{\gamma}^{T_9} - tk^{T_9}\right\}.$$

Proceeding in the same fashion for the remaining T_i then gives the $\mathbb{P}^{T_{10}}$-dynamics

$$L(t,T_i) = L(0,T_i)\exp\left\{\tilde{L}_t^{T_{10}} + t\bar{\gamma}^{T_{i+1}} - tk^{T_{i+1}}\right\}.$$

The numerical values of $\bar{\gamma}^{T_i}$ and k^{T_i} can be found in Table 4.1.

Now that we have worked out all parameters, we can turn to the simulation procedure, which turns out to be remarkably simple. Assume that the interest rate derivative to price

T_i	$\bar{\gamma}^{T_i}$	k^{T_i}
10	-0.013937	0
9	-0.013961	0.001332
8	-0.013985	0.002666
7	-0.014012	0.004004
6	-0.014040	0.005344
5	-0.014070	0.006688
4	-0.014101	0.008035
3	-0.014134	0.009386
2	-0.014169	0.010740

Table 4.1: Drift terms

depends only the LIBORs observed at times T_i (and not on their intermediate values).[18] Simulation can be performed by the following algorithm:[19]

For $i = 1$ to $n - 1$

Generate a $VG(C, G)$-distributed random number R_i.

For $j = i$ to $n - 1$

$$L(T_i, T_j) := L(T_{i-1}, T_j) \exp\left\{R_i + \bar{\gamma}^{T_j+1} - k^{T_j+1}\right\}.$$

The result of the above algorithm will be an upper diagonal matrix that represents the evolution of the whole LIBOR curve through time. This forms the basis for pricing interest rate derivatives.

4.5.2 Implementation of the exact model

Implementation of the exact model is more cumbersome, the main reason being the stochasticity of the compensators. After each step of the simulation procedure, the compensators have to be calculated anew. This brings about the problem of having to re-evaluate

[18]This is a very weak restriction that most interest rate derivatives obey. If necessary, it can be easily relaxed.

[19]We assume the $L(0, T_i)$ to be properly initialised.

the integrals involving the compensators after each step of the simulation, as can be seen from formulae (4.4) and (4.5). In general, one cannot expect closed formulae for these integrals, and consequently one would have to perform several numerical integrations after each timestep, rendering the simulation extremely slow and almost infeasible. For our concrete parameterisation, however, closed formulae for the relevant integrals exist, with the consequence that simulation of the exact model is practicable and not prohibitively time consuming. We will take advantage of this in the next section.

4.6 Testing our approximation

In this section, we test the impact of the approximation proposed on zero-bond, caplet and swaption prices. As already mentioned, in our concrete setup, it is possible to simulate not only the approximate, but also the exact model efficiently, which allows us to calculate the prices of the aforementioned products based on 10 million antithetic paths, thereby keeping the simulation error very low. A great number of paths is of particular relevance in this context, as we want to be able to gauge the error due to our approximation while keeping the distortion due to the simulation error[20] as low as possible. We choose a size of 1 year for the timesteps, and thus incur a certain discretization error. It is this discretization error the pricing errors generated by the exact LMM stem largely from, while the simulation error, as measured by the standard error, is almost negligible. In contrast, the approximate model does *not* suffer from discretization error, i.e. decreasing the step-size of the simulation to values less than 1 would not give more accurate results.[21]

Table 4.2 contrasts the prices of zero bonds with face value 1 calculated with the exact (column exactLMM) and approximate (column proxLMM) models with the exact bond prices. Both (relative) differences and standard errors for both models are of negligible

[20] Here, by simulation error we mean the deviation of the calculated arithmetic mean from the 'true' expected value. It is not to be confused with the discretization error that stems from approximating a continuous time model by a discrete one.

[21] This property of the approximate LMM is rather a peculiarity of our concrete parameterisations than a general fact. Were we to choose e.g. functions λ that change during the timesteps of the simulation, it would no longer hold true.

Maturity	exactLMM			proxLMM			Exact
T_i	Mean	Std. err.	Err.	Mean	Std. err.	Err.	
1	0.95241	0.00003	0.003%	0.95241	0.00003	0.003%	0.95238
2	0.90707	0.00004	0.005%	0.90711	0.00003	0.009%	0.90703
3	0.86388	0.00004	0.005%	0.86399	0.00005	0.017%	0.86384
4	0.82273	0.00004	0.004%	0.82287	0.00004	0.021%	0.82270
5	0.78354	0.00004	0.002%	0.78368	0.00003	0.020%	0.78353
6	0.74621	0.00004	-0.001%	0.74633	0.00002	0.015%	0.74622
7	0.71066	0.00002	-0.002%	0.71074	0.00002	0.008%	0.71068
8	0.67682	0.00002	-0.003%	0.67684	0.00001	0.001%	0.67684
9	0.64460	0.00001	-0.002%	0.64459	0.00001	-0.003%	0.64461

Table 4.2: Bond prices

magnitude.

The situation looks different for caplets, where we compare the prices of at-the-money caplets that pay $10000(L(T_i,T_i)-0.05)^+$ at time T_{i+1}. The exact prices reported in Table 4.3 are obtained by integrating the caplet payoff against the probability density of the corresponding LIBOR-realisation. In the case of the exact LMM, comparing the pricing errors with the corresponding standard errors, one can conclude that the former must largely stem from discretization, i.e. they can be reduced by choosing smaller timesteps. Even for a timestep as large as one year, the discretization error is well within tolerable limits. For the approximate LMM, the maximum relative error is roughly twice as large, but still reasonably small. In both cases, the errors become smaller for maturities approaching 9. This observation conforms with the intuition that the approximation errors get larger the further we move away from the terminal measure $\mathbb{P}^{T_{10}}$ and the associated LIBOR process $(L(t,T_9))$, which – both for the exact and approximate LMMs – suffers neither from discretization nor approximation error.

Next, we consider the prices of at-the-money 5-into-5 payer swaptions (see Table 4.4). In absence of an exact swaption price, we can only compare the two model-prices directly. The picture is very much the same as before: The prices are almost identical, the relative

Maturity	exactLMM			proxLMM			Exact
T_i	Mean	Std. err.	Err.	Mean	Std. err.	Err.	
1	28.83	0.07	1.49 %	28.88	0.04	1.65 %	28.41
2	39.94	0.08	1.36 %	40.25	0.06	2.15 %	39.41
3	46.97	0.09	1.26 %	47.60	0.08	2.61 %	46.39
4	51.76	0.10	1.06 %	52.69	0.08	2.87 %	51.22
5	55.10	0.13	0.84 %	56.20	0.11	2.84 %	54.64
6	57.39	0.12	0.58 %	58.50	0.13	2.53 %	57.06
7	58.89	0.10	0.30 %	59.83	0.13	1.90 %	58.72
8	59.81	0.09	0.06 %	60.39	0.14	1.02 %	59.78
9	60.27	0.08	-0.16 %	60.34	0.15	-0.05 %	60.37

Table 4.3: Caplet prices

exactLMM		proxLMM	
Mean	Standard error	Mean	Standard error
249.62	0.48	251.69	0.29

Table 4.4: 5x5 swaption prices in basis points

difference being less than one percent.

To conclude, our simulation experiments show that the approximate LMM is both an efficient and accurate alternative to the exact LMM for the model at hand. In case the concrete parameterisation of the exact LMM does not admit closed formulae for the integrals involving compensators, the approximate LMM is the only viable alternative. Our simulation experiments should be extended to a broader range of driving processes and parameterisations to further substantiate our findings, but for reasons of scope, we leave that for future research.

4.7 Implied volatilities and their dynamics in a Lévy LMM

As already pointed out in Chapter 2, a paramount criterion for the adequacy of a modelling approach is not only its ability to statically reproduce observed implied volatility patterns, but also the dynamics of implied volatilities it induces. In this section, we investigate both the static and dynamic properties of implied volatilities induced by Lévy LMMs. We commence with the dynamic point of view.

4.7.1 Smile dynamics

The following two propositions contrast the dynamic behaviour of implied volatility surfaces in time-homogeneous versus time-inhomogeneous exponential Lévy models.

Proposition 4.1 (Smile dynamics in a time-homogeneous exponential Lévy model). *In a time-homogeneous exponential Lévy model, Black implied volatilities of caplets are a function of moneyness $m = L/K$ and time to maturity $\tau = T - t$ only. In other words, time-homogeneous exponential Lévy models exhibit the so-called* forward-propagated smile property.[22]

[22] Compare Rebonato [2004], p.593.

Proof: Denote the price of a caplet with strike K, maturing in T and paying in $T+\delta$ by $C^{\text{Lévy}}(L,K,t,T)$, where t is the current time and $L = L(t,T)$. Then

$$\begin{aligned}
C^{\text{Lévy}}(L,K,t,T) &= p(t,T+\delta)\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[(L(T,T)-K)^+|L(t,T)=L\right] \\
&= p(t,T+\delta)\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[\left(L(0,T)e^{L_t}e^{L_T-L_t}-K\right)^+|L(t,T)=L\right] \\
&= p(t,T+\delta)\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[\left(Le^{L_\tau}-K\right)^+\right] \\
&= p(t,T+\delta)K\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[\left(\frac{L}{K}e^{L_\tau}-1\right)^+\right]
\end{aligned}$$

and thus

$$\begin{aligned}
\frac{C^{\text{Lévy}}(L,K,t,T)}{L} &= \frac{p(t,T+\delta)}{m}\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[\left(me^{L_\tau}-1\right)^+\right] \\
&= p(t,T+\delta)g^{\text{Lévy}}(m,\tau).
\end{aligned}$$

As the Black-Scholes model is an exponential Lévy model, we get

$$\frac{C^{\text{BS}}(L,K,t,T,\sigma)}{L} = p(t,T+\delta)g^{\text{BS}}(m,\tau,\sigma).$$

We conclude that the implied volatility function σ^{imp}, which is implicitly defined as the solution to the equation

$$g^{\text{Lévy}}(m,\tau) = g^{\text{BS}}(m,\tau,\sigma^{\text{imp}}),$$

must have the form $\sigma^{\text{imp}}(m,\tau)$. □

Proposition 4.2 (Smile dynamics in a time-inhomogeneous exponential Lévy model). *In a time-inhomogeneous exponential Lévy model, Black implied volatilities of caplets are a function of moneyness $m = L/K$, current time t and maturity T only. In the terminology of Rebonato [2004], time-inhomogeneous exponential Lévy models exhibit the so-called* floating smile property.[23]

[23] Other authors, e.g. Derman [1999], call this the *sticky moneyness property.*

Proof:

$$\begin{aligned} C^{\text{Lévy}}(L,K,t,T) &= p(t,T+\delta)\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[(L(T,T)-K)^+|L(t,T)=L\right] \\ &= p(t,T+\delta)\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[\left(L(0,T)e^{L_t}e^{L_T-L_t}-K\right)^+|L(t,T)=L\right] \\ &= p(t,T+\delta)\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[\left(Le^{L_T-L_t}-K\right)^+\right] \end{aligned}$$

and thus

$$\begin{aligned} \frac{C^{\text{Lévy}}(L,K,t,T)}{L} &= \frac{p(t,T+\delta)}{m}\mathbb{E}_{\mathbb{P}^{T+\delta}}\left[\left(me^{L_T-L_t}-1\right)^+\right] \\ &= p(t,T+\delta)g^{\text{Lévy}}(m,t,T). \end{aligned}$$

For the Black-Scholes model, we get

$$\frac{C^{\text{BS}}(L,K,t,T,\sigma)}{L} = p(t,T+\delta)g^{\text{BS}}(m,t,T,\sigma).$$

We conclude that the implied volatility function σ^{imp}, which is implicitly defined as the solution to the equation

$$g^{\text{Lévy}}(m,t,T) = g^{\text{BS}}(m,t,T,\sigma^{\text{imp}})$$

must have the form $\sigma^{\text{imp}}(m,t,T)$. □

It is obvious how to interpret the above results: As is the case for local volatility models, the future implied volatility surface (observed in terms of moneyness) for any future date t is already known today.[24] In a time-homogeneous model, the implied volatility surface, when observed in terms of moneyness and time to maturity, remains constant, which means that neither moves in the underlying nor the passage of time alter the implied volatility surface: it remains stationary in time to maturity and moneyness. In a time-inhomogeneous model, however, the passage of time changes the shape of the implied volatility surface, while moves in the underlying (with t held fixed) don't. Thus, while time-inhomogeneous models give us much more flexibility when it comes to fitting term structures, they bring about non-stationarity of implied volatility surfaces, which is generally deemed a rather undesirable property. We will elaborate on this shortly.

[24] But observe that for local volatility models, the knowledge of the t-implied volatility surface is conditional on knowing the realisation of the underlying at time t.

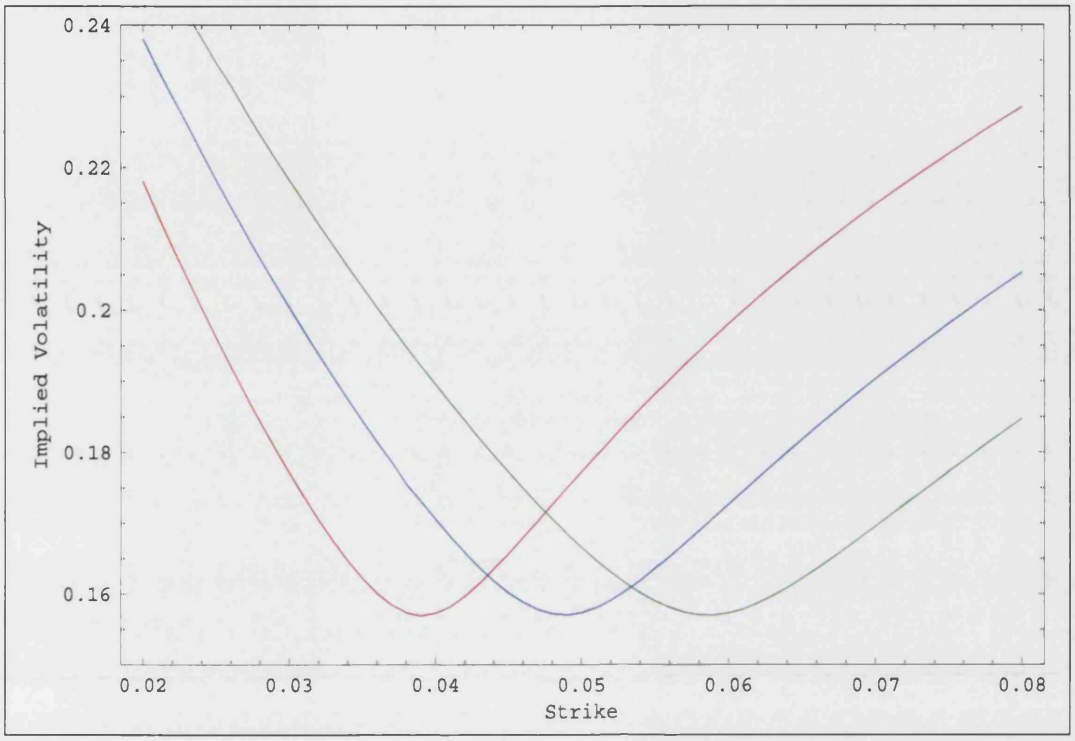

Figure 4.1: Variance Gamma one year caplet smiles in terms of strike

We shall now visualise the caplet smile dynamics in our Variance Gamma-model, and contrast them with the corresponding dynamics obtained in a local volatility model. In both cases, we will look at smiles of caplets on $L(1,1)$ paying in $T = 2$. We confine ourselves to smiles (rather than whole surfaces), because qualitatively, the corresponding surface dynamics are exactly the same.

Figure 4.1 shows Variance Gamma smiles in terms of absolute strike levels for $L(0,1) = 0.04$ (red line), $L(0,1) = 0.05$ (blue line) and $L(0,1) = 0.06$ (green line). Apparently, the smiles move in the same direction as the underlying. In stark contrast, one observes in Figure 4.2 that in local volatility models, smiles move in the *opposite* direction of the underlying. As already mentioned in Chapter 2, Hagan *et al.* [2002] remark that this contradicts real-world smile dynamics, where smiles move in the same direction as the underlying. Thus, smile dynamics induced by Lévy models are considerably more realistic than those induced by their local volatility counterparts.[25]

[25] Of course, Lévy smile dynamics are also only a crude approximation to reality. As documented by e.g. Derman [1999] or Cont & da Fonseca [2002], empirical implied volatilities quoted in terms of moneyness

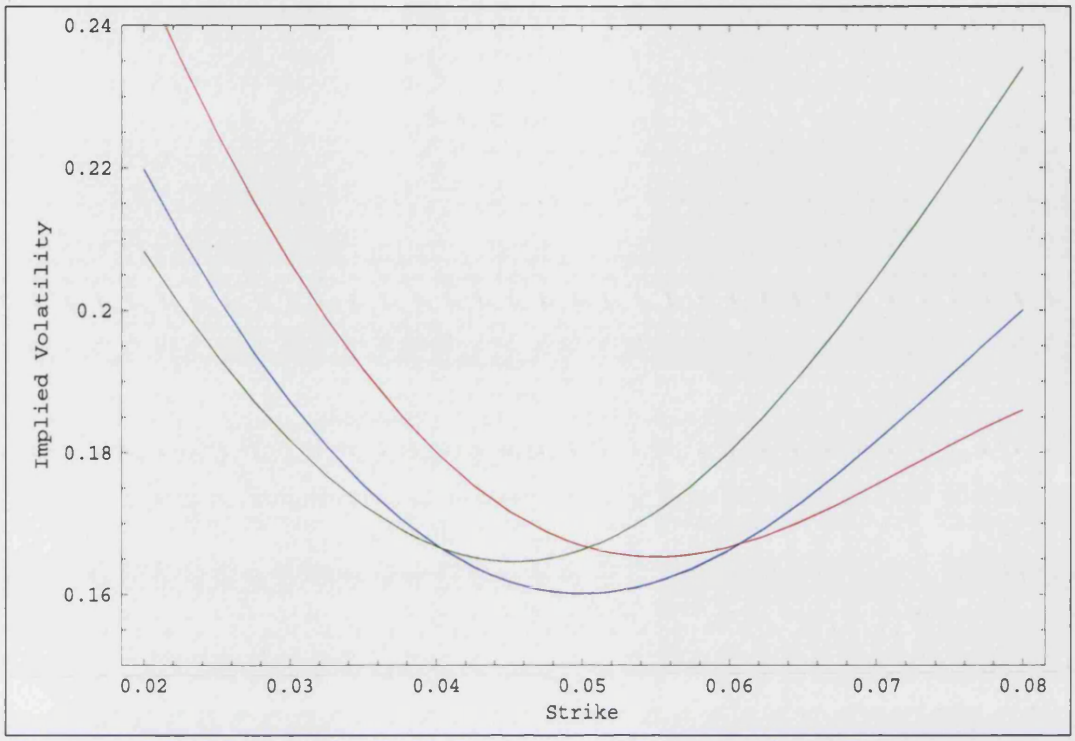

Figure 4.2: Local volatility one year caplet smiles in terms of strike

Figures 4.3 and 4.4 display the smile dynamics in terms of moneyness. While Figure 4.3 shows the sticky moneyness property, Figure 4.4 makes the characteristics of local volatility models already observed for absolute strikes even more obvious.

4.7.2 Implied volatility surfaces

Now we consider the qualitative properties of implied volatility surfaces induced by a time-homogeneous Lévy model on the basis of our Variance Gamma example.

As Figure 4.5 shows, our model produces a realistic-looking implied volatility surface for short maturities. For medium to long maturities, however, the smile becomes almost flat, while typical real-world caplet implied volatility surfaces show considerably more pronounced smile patterns (see Jarrow *et al.* [2003]). This so-called flattening-out effect is a consequence of the central limit theorem.[26] Similarly, when the model is calibrated

show considerable variability, which contradicts the sticky smile property.

[26] But see also Carr & Wu [2003] for an exponential Lévy model based on an α-stable Lévy process with

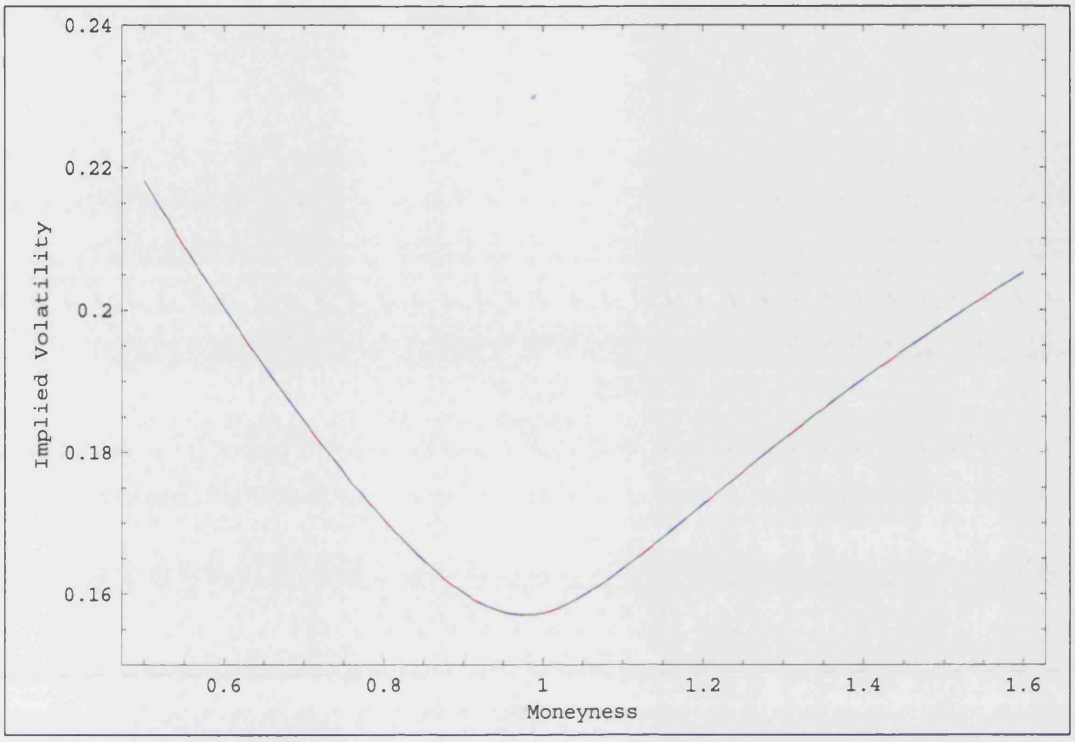

Figure 4.3: Variance Gamma one year caplet smiles in terms of moneyness

to the smile-patterns at the long end of the maturity spectrum, smiles at the short end will typically be much too steep to be consistent with real-world ones. This dilemma can be avoided by using time-inhomogeneous Lévy models, which allow simultaneous and almost perfect fits to all observed caplet smiles. But this often comes at the price of a high degree of non-stationarity: In order to match real-world implied volatility patterns, the parameters typically have to be chosen in a way that makes these models highly non-stationary. This property leads to future implied volatility surfaces that are very different from today's, which is undesirable, one of the reasons being that it gives rise to unrealistic prices of derivatives that strongly depend on future implied volatility surfaces, such as forward start options or cliquets.

For local volatility models, there is a phenomenon that is related to the just described flattening-out effect. Local volatility models, when calibrated to implied volatility surfaces

maximum negative skewness, where the central limit theorem does not apply and the implied volatility surface does not flatten out.

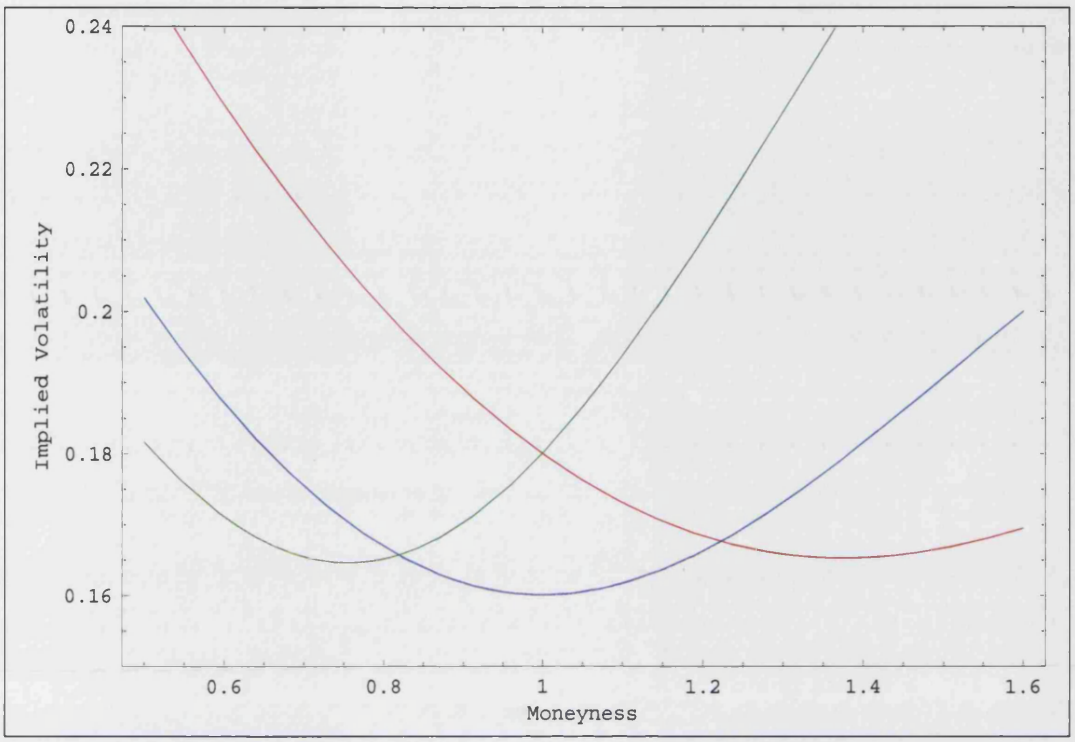

Figure 4.4: Local volatility one year caplet smiles in terms of moneyness

that are steep for short maturities and relatively flat for longer ones, predict future implied volatility surfaces that are 'flatter' than the current one. The reason is clear from the discussion in Chapter 2: In the course of time, the 'short end' of the local volatility function that generates the steep patterns for shorter maturities 'disappears', i.e. loses its impact on the implied volatility surface. Thus, flattening out in local volatility models is a dynamic property that depends on the initial shape of the implied volatility surface, while in time-homogenous Lévy models, it is a static one.

Summarising, we can say that in terms of smile (respectively surface) dynamics, Lévy models are clearly superior to their local volatility counterparts, as for the former, smiles (or, more generally speaking, implied volatility structures when observed in terms of absolute strike levels) move in the same direction as the underlying, while for the latter, the opposite is the case. The cross-sectional performance (that is, the capability of reproducing a smile for a certain maturity) of both local volatility and Lévy models is excellent. However, when it comes to fitting whole volatility surfaces, time-homogeneous Lévy mod-

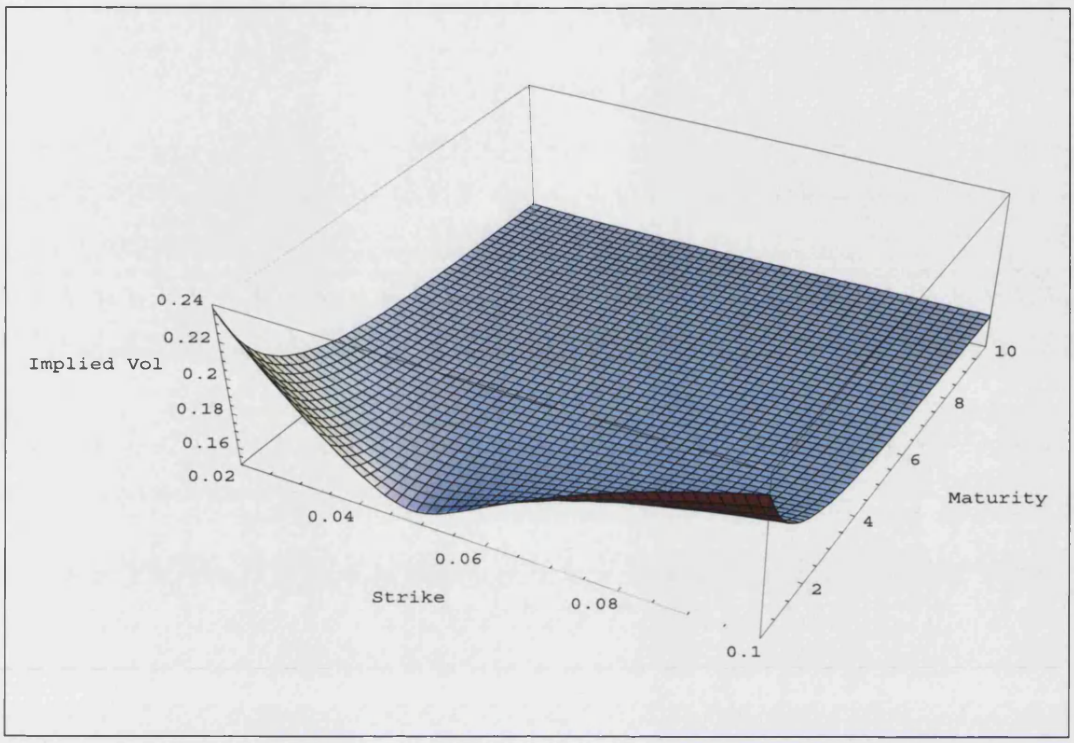

Figure 4.5: Variance Gamma caplet implied volatility surface

els are severely limited, while time-inhomogeneous Lévy and local volatility models are not, but this flexibility is bought dearly by a typically high degree of non-stationarity associated with the unwanted side-effects outlined above. This leads us to conclude that, among the three model classes considered, time-inhomogeneous Lévy models are the class of choice when the emphasis is on an accurate fit to observed implied volatility surfaces and plausible smile dynamics, and the derivative to price is not overly sensitive to the non-stationarity exhibited by this model class.

A discussion on Lévy models is not complete without a few words on market incompleteness, the choice of an equivalent martingale measure, calibration and hedging. Apart from few exceptions (when the driving noise is a Brownian motion or a Poisson process), Lévy models are incomplete, and therefore there is no unique martingale measure. But, as we follow an implied approach, this does not pose a problem: If we choose the driving Lévy process from of a certain parametric class, we can obtain its parameters by calibrating the model prices (respectively model implied volatilities) to those observed in the market,

which basically means that we let the market choose the measure.[27]

For reasons of scope, we do not address the issue of hedging in Lévy models here. Suffice it to say that, as perfect hedging in incomplete markets is (by definition of an incomplete market) not possible for all contingent claims, one has to resort to approximate hedging strategies (for instance minimum-variance hedging). Schoutens [2003] and Cont & Tankov [2004] provide overviews.

4.8 Summary

In this chapter, we give a novel derivation of LIBOR dynamics and measure relationships in a Lévy LMM using a change of numéraire argument. We develop an approximation technique that simplifies and speeds up implementation and simulation considerably. Subsequently, we test our approximation on the basis of a concrete numerical example, and find that it performs well. We discuss numerical issues involved in the implementation of a Lévy LMM and point out possible complications and limitations. Finally, we explore the smile dynamics induced by our Lévy LMM and contrast them with those encountered in LMMs based on local volatility functions.

[27] Calibration techniques for Lévy models can be found in Cont & Tankov [2004], p.463.

Chapter 5

Summary and conclusion

The objective of this dissertation is to develop smile-consistent financial models with a special emphasis on LIBOR market models.

Chapter 1 provides a general overview of approaches to smile modelling. In Chapter 2, we propose analytical approximate solutions to the single smile problem. Extensive numerical tests based on several extreme volatility scenarios show that our methods provide an excellent fit to the input data. The methods we detail in this chapter are tailor-made for the construction of smile-consistent LIBOR market models. In Chapter 3, we develop the theory of generalised extended LMMs. Relying on our results from Chapter 2, we are able to fit any given set of implied volatility smiles for the maturities of the tenor structure (almost) exactly, while at the same time preserving numerical tractability. Moreover, we propose a swaption approximation, which we subject to numerical testing, and find that its quality in the test-cases considered is excellent. In Chapter 4, we give a novel derivation of LIBOR dynamics and measure relationships in a Lévy-driven LMM using a change of numéraire argument. We discuss an approximation technique that simplifies and speeds up implementation and simulation considerably. Subsequently, we test our approximation on a Lévy LMM driven by a Variance Gamma process, and find that it performs well. We discuss numerical issues involved in the implementation of Lévy LMMs and point out possible problems and limitations. Finally, we contrast the implied volatility dynamics induced by a Lévy LMM with those encountered in LMMs based on

local volatility functions, and find that the former are more realistic.

Both model classes considered – GEMMs and Lévy LMMs – are capable of providing an excellent fit to given implied volatility surfaces. However, they differ with regard to numerical tractability, induced smile dynamics, and market completeness. While GEMMs are easier to handle and preserve completeness, Lévy LMMs feature more realistic smile dynamics. However, both GEMMs and time-inhomogeneous Lévy LMMs suffer from non-stationarity. Thus, while we can consider the problem of smile-consistent modelling in a LIBOR context solved, a challenging problem remains: To identify a class of LIBOR market models that is flexible enough to fit real-world implied volatility surfaces, while at the same time giving rise to realistic implied volatility dynamics and preserving stationarity. A possible way to tackle this issue could be to develop Lévy-based LMMs that also incorporate stochastic volatility, e.g. by extending the approach pursued in Carr *et al.* [2003] to the interest rate world. We are confident that this approach can improve on the current state of the art, but we have to leave the proof for future research.

Appendix A

Numerical Results

Strike	Exact price	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
10	90.000	0.000	0.000	0.000	0.000	0.000	0.000
20	80.001	0.000	0.000	0.000	0.000	0.000	0.000
30	70.004	0.000	0.000	0.000	0.000	0.000	0.000
40	60.011	0.000	-0.001	-0.001	-0.001	-0.001	-0.001
50	50.031	0.001	-0.001	-0.001	-0.001	-0.001	-0.001
60	40.086	0.003	-0.001	-0.001	-0.001	-0.003	-0.003
70	30.249	0.007	-0.003	-0.003	-0.003	-0.005	-0.005
80	20.738	0.013	-0.006	-0.006	-0.006	-0.008	-0.008
90	12.116	0.017	-0.010	-0.010	-0.010	-0.010	-0.010
100	5.416	0.016	-0.013	-0.013	-0.013	-0.010	-0.010
110	1.579	0.009	-0.009	-0.009	-0.009	-0.008	-0.008
120	0.245	0.002	-0.003	-0.003	-0.003	-0.003	-0.003
130	0.016	0.000	0.000	0.000	0.000	0.000	0.000
140	0.000	0.000	0.000	0.000	0.000	0.000	0.000
150	0.000	0.000	0.000	0.000	0.000	0.000	0.000
160	0.000	0.000	0.000	0.000	0.000	0.000	0.000
170	0.000	0.000	0.000	0.000	0.000	0.000	0.000
180	0.000	0.000	0.000	0.000	0.000	0.000	0.000
190	0.000	0.000	0.000	0.000	0.000	0.000	0.000
200	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table A.1: Approximation errors for the 1 month skew case. Δ P is the price difference between the model price and the exact price. Positive figures indicate overpricing by the model.

Strike	Exact price	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
10	90.000	0.000	0.000	0.000	0.000	0.000	0.000
20	80.000	0.000	0.000	0.000	0.000	0.000	0.000
30	70.000	0.000	0.000	0.000	0.000	0.000	0.000
40	60.000	0.000	0.000	0.000	0.000	0.000	0.000
50	50.001	0.000	0.000	0.000	0.000	0.000	0.000
60	40.003	0.001	0.000	0.000	0.000	0.000	0.000
70	30.019	0.003	0.000	0.000	0.000	-0.001	-0.001
80	20.147	0.012	-0.001	-0.001	-0.001	-0.003	-0.003
90	11.028	0.037	-0.002	-0.002	-0.002	-0.003	-0.003
100	4.426	0.059	-0.001	0.000	0.000	0.006	0.006
110	1.321	0.048	0.002	0.002	0.002	0.001	0.001
120	0.350	0.026	0.002	0.002	0.002	-0.002	-0.002
130	0.101	0.013	0.002	0.002	0.002	-0.001	-0.001
140	0.035	0.007	0.001	0.001	0.001	-0.001	-0.001
150	0.015	0.004	0.001	0.001	0.001	-0.001	-0.001
160	0.008	0.003	0.000	0.000	0.000	0.000	0.000
170	0.005	0.002	0.000	0.000	0.000	0.000	0.000
180	0.003	0.001	0.000	0.000	0.000	0.000	0.000
190	0.003	0.001	0.000	0.000	0.000	0.000	0.000
200	0.002	0.001	0.000	0.000	0.000	0.000	0.000

Table A.2: Approximation errors for the 1 month smile case. Δ P is the price difference between the model price and the exact price. Positive figures indicate overpricing by the model.

Strike	Exact price	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
10	90.000	0.000	0.000	0.000	0.000	0.000	0.000
20	80.007	-0.001	0.000	0.000	0.000	0.000	0.000
30	70.067	-0.005	-0.002	-0.002	-0.002	0.000	0.000
40	60.276	-0.015	-0.005	-0.005	-0.005	-0.002	-0.001
50	50.774	-0.027	-0.010	-0.009	-0.009	-0.007	-0.003
60	41.732	-0.036	-0.014	-0.013	-0.013	-0.013	-0.007
70	33.341	-0.040	-0.016	-0.016	-0.016	-0.019	-0.011
80	25.789	-0.037	-0.017	-0.016	-0.016	-0.023	- 0.013
90	19.234	-0.028	-0.015	-0.014	-0.014	-0.025	-0.013
100	13.780	-0.017	-0.012	-0.012	-0.012	-0.023	-0.012
110	9.450	-0.005	-0.009	-0.008	-0.008	-0.020	-0.009
120	6.184	0.003	-0.006	-0.005	-0.005	-0.016	-0.007
130	3.854	0.008	-0.003	-0.003	-0.003	-0.013	-0.005
140	2.283	0.010	-0.002	-0.002	-0.002	-0.010	-0.004
150	1.284	0.009	-0.001	-0.001	-0.001	-0.007	-0.003
160	0.687	0.007	0.000	0.000	0.000	-0.005	-0.003
170	0.349	0.005	0.000	0.000	0.000	-0.003	-0.002
180	0.169	0.003	0.000	0.000	0.000	-0.002	-0.001
190	0.078	0.002	0.000	0.000	0.000	-0.001	-0.001
200	0.035	0.001	0.000	0.000	0.000	-0.001	0.000
210	0.015	0.001	0.000	0.000	0.000	0.000	0.000
220	0.006	0.000	0.000	0.000	0.000	0.000	0.000
230	0.002	0.000	0.000	0.000	0.000	0.000	0.000
240	0.001	0.000	0.000	0.000	0.000	0.000	0.000
250	0.000	0.000	0.000	0.000	0.000	0.000	0.000
260	0.000	0.000	0.000	0.000	0.000	0.000	0.000
270	0.000	0.000	0.000	0.000	0.000	0.000	0.000
280	0.000	0.000	0.000	0.000	0.000	0.000	0.000
290	0.000	0.000	0.000	0.000	0.000	0.000	0.000
300	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table A.3: Approximation errors for the 1 year skew case. Δ P is the price difference between the model price and the exact price. Positive figures indicate overpricing by the model.

Strike	Exact price	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
10	90.000	0.000	0.000	0.000	0.000	0.000	0.000
20	80.002	-0.001	-0.001	-0.001	-0.001	0.000	0.000
30	70.020	-0.002	-0.004	-0.004	-0.003	-0.002	-0.002
40	60.066	0.004	-0.007	-0.006	-0.006	-0.006	-0.006
50	50.155	0.026	-0.006	-0.005	-0.005	-0.011	-0.011
60	40.354	0.066	-0.004	-0.003	-0.002	-0.015	-0.015
70	30.852	0.121	-0.004	-0.002	-0.001	-0.017	-0.017
80	22.063	0.190	-0.006	-0.003	-0.002	-0.010	-0.011
90	14.629	0.257	-0.008	-0.004	-0.002	0.010	0.009
100	9.083	0.295	-0.006	-0.002	-0.001	0.027	0.025
110	5.447	0.294	-0.002	0.002	0.003	0.018	0.017
120	3.282	0.267	0.002	0.006	0.007	0.003	0.002
130	2.052	0.231	0.005	0.008	0.009	-0.007	-0.008
140	1.355	0.195	0.006	0.008	0.009	-0.013	-0.014
150	0.949	0.163	0.005	0.007	0.008	-0.015	-0.016
160	0.703	0.135	0.003	0.005	0.006	-0.016	-0.017
170	0.545	0.112	0.001	0.003	0.004	-0.016	-0.017
180	0.439	0.091	-0.001	0.000	0.001	-0.016	-0.016
190	0.365	0.074	-0.004	-0.002	-0.002	-0.016	-0.016
200	0.310	0.059	-0.006	-0.005	-0.004	-0.015	-0.016
210	0.268	0.046	-0.008	-0.007	-0.007	-0.015	-0.015
220	0.234	0.036	-0.010	-0.009	-0.008	-0.014	-0.015
230	0.207	0.027	-0.011	-0.010	-0.010	-0.014	-0.014
240	0.184	0.019	-0.012	-0.011	-0.011	-0.013	-0.013
250	0.165	0.013	-0.013	-0.012	-0.012	-0.013	-0.013
260	0.148	0.008	-0.014	-0.013	-0.012	-0.012	-0.012
270	0.133	0.004	-0.014	-0.013	-0.012	-0.011	-0.011
280	0.119	0.001	-0.014	-0.013	-0.012	-0.011	-0.011
290	0.107	-0.002	-0.013	-0.013	-0.012	-0.010	-0.010
300	0.097	-0.003	-0.013	-0.012	-0.012	-0.009	-0.009

Table A.4: Approximation errors for the 1 year smile case. Δ P is the price difference between the model price and the exact price. Positive figures indicate overpricing by the model.

Strike	Exact price	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
10	90.060	-0.003	0.001	0.001	0.001	0.001	0.001
20	80.610	-0.028	0.001	0.002	0.002	0.005	0.005
30	71.967	-0.074	-0.001	0.000	0.000	0.005	0.005
40	64.179	-0.127	-0.005	-0.003	-0.003	-0.002	-0.002
50	57.200	-0.177	-0.009	-0.007	-0.007	-0.012	-0.012
60	50.959	-0.220	-0.013	-0.011	-0.010	-0.024	-0.025
70	45.382	-0.252	-0.015	-0.012	-0.012	-0.036	-0.036
80	40.400	-0.272	-0.015	-0.012	-0.012	-0.044	-0.045
90	35.949	-0.281	-0.014	-0.010	-0.010	-0.049	-0.050
100	31.975	-0.281	-0.010	-0.007	-0.007	-0.050	-0.050
110	28.428	-0.274	-0.006	-0.003	-0.003	-0.046	-0.047
120	25.264	-0.261	-0.002	0.002	0.002	-0.040	-0.041
130	22.442	-0.244	0.004	0.007	0.007	-0.033	-0.033
140	19.928	-0.224	0.009	0.012	0.013	-0.025	-0.025
150	17.689	-0.202	0.014	0.017	0.018	-0.017	-0.017
160	15.697	-0.179	0.019	0.022	0.022	-0.009	-0.010
170	13.926	-0.157	0.023	0.026	0.027	-0.003	-0.003
180	12.352	-0.135	0.027	0.030	0.030	0.003	0.002
190	10.954	-0.115	0.031	0.033	0.033	0.007	0.007
200	9.714	-0.095	0.033	0.036	0.036	0.011	0.011
210	8.615	-0.078	0.035	0.037	0.038	0.014	0.014
220	7.640	-0.062	0.037	0.039	0.039	0.016	0.016
230	6.776	-0.048	0.038	0.040	0.040	0.018	0.017
240	6.011	-0.035	0.038	0.040	0.040	0.018	0.018
250	5.334	-0.024	0.038	0.040	0.040	0.019	0.019
260	4.735	-0.014	0.038	0.040	0.040	0.019	0.019
270	4.204	-0.006	0.038	0.039	0.039	0.019	0.018
280	3.735	0.001	0.037	0.038	0.038	0.018	0.018
290	3.320	0.006	0.036	0.037	0.037	0.017	0.017
300	2.952	0.011	0.035	0.036	0.036	0.016	0.016

Table A.5: Approximation errors for the 10 year skew case. Δ P is the price difference between the model price and the exact price. Positive figures indicate overpricing by the model.

Strike	Exact price	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
10	90.012	-0.002	-0.002	-0.002	-0.002	-0.001	-0.001
20	80.224	-0.022	-0.029	-0.028	-0.027	-0.014	-0.014
30	70.841	-0.022	-0.079	-0.075	-0.073	-0.044	-0.045
40	61.873	0.075	-0.110	-0.103	-0.100	-0.065	-0.067
50	53.373	0.274	-0.111	-0.100	-0.096	-0.060	-0.063
60	45.484	0.531	-0.094	-0.078	-0.073	-0.030	-0.034
70	38.376	0.790	-0.074	-0.053	-0.046	0.019	0.013
80	32.183	1.013	-0.058	-0.032	-0.023	0.075	0.067
90	26.954	1.176	-0.049	-0.018	-0.008	0.123	0.113
100	22.653	1.273	-0.046	-0.012	-0.002	0.148	0.138
110	19.178	1.310	-0.048	-0.014	-0.003	0.144	0.133
120	16.400	1.297	-0.056	-0.022	-0.011	0.120	0.109
130	14.182	1.246	-0.069	-0.036	-0.025	0.088	0.077
140	12.407	1.169	-0.085	-0.053	-0.043	0.053	0.043
150	10.973	1.075	-0.104	-0.074	-0.064	0.019	0.010
160	9.803	0.972	-0.124	-0.096	-0.087	-0.012	-0.020
170	8.836	0.866	-0.145	-0.119	-0.110	-0.039	-0.047
180	8.026	0.761	-0.166	-0.141	-0.133	-0.063	-0.070
190	7.338	0.660	-0.185	-0.161	-0.154	-0.083	-0.090
200	6.745	0.565	-0.203	-0.180	-0.173	-0.100	-0.107
210	6.230	0.477	-0.219	-0.197	-0.190	-0.114	-0.120
220	5.776	0.396	-0.232	-0.211	-0.205	-0.126	-0.131
230	5.372	0.323	-0.244	-0.223	-0.217	-0.134	-0.140
240	5.010	0.258	-0.253	-0.233	-0.227	-0.141	-0.146
250	4.682	0.200	-0.260	-0.241	-0.235	-0.145	-0.150
260	4.384	0.149	-0.265	-0.246	-0.241	-0.148	-0.153
270	4.111	0.104	-0.268	-0.250	-0.244	-0.150	-0.154
280	3.861	0.065	-0.269	-0.252	-0.246	-0.150	-0.154
290	3.629	0.032	-0.269	-0.252	-0.247	-0.149	-0.153
300	3.415	0.003	-0.267	-0.251	-0.246	-0.147	-0.151

Table A.6: Approximation errors for the 10 year smile case. Δ P is the price difference between the model price and the exact price. Positive figures indicate overpricing by the model.

Strike	Exact imp vol	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
50	93.58%	0.26%	-0.34%	-0.31%	-0.31%	-0.48%	-0.48%
60	80.24%	0.39%	-0.19%	-0.19%	-0.19%	-0.36%	-0.36%
70	69.32%	0.36%	-0.15%	-0.15%	-0.15%	-0.26%	-0.26%
80	60.38%	0.28%	-0.13%	-0.13%	-0.13%	-0.18%	-0.18%
90	53.06%	0.20%	-0.12%	-0.12%	-0.12%	-0.12%	-0.12%
100	47.07%	0.14%	-0.11%	-0.11%	-0.11%	-0.08%	-0.08%
110	42.16%	0.10%	-0.10%	-0.10%	-0.10%	-0.09%	-0.09%
120	38.14%	0.07%	-0.09%	-0.09%	-0.09%	-0.11%	-0.11%
130	34.85%	0.05%	-0.08%	-0.08%	-0.08%	-0.11%	-0.11%
140	32.16%	0.04%	-0.07%	-0.07%	-0.07%	-0.10%	-0.10%
150	29.96%	0.04%	-0.06%	-0.06%	-0.06%	-0.10%	-0.10%

Table A.7: Approximation errors for the 1 month skew case. Δ P is the difference between the model implied volatility and the exact one. Positive figures indicate overpricing by the model.

Strike	Exact imp vol	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
50	69.56%	1.28%	0.00%	0.00%	0.00%	-0.50%	-0.50%
60	57.18%	1.25%	0.04%	0.04%	0.04%	-0.53%	-0.53%
70	48.19%	0.85%	-0.02%	0.01%	0.01%	-0.31%	-0.31%
80	42.41%	0.62%	-0.03%	-0.03%	-0.03%	-0.17%	-0.17%
90	39.35%	0.52%	-0.03%	-0.02%	-0.02%	-0.04%	-0.04%
100	38.45%	0.52%	-0.01%	0.00%	0.00%	0.05%	0.05%
110	39.19%	0.57%	0.02%	0.03%	0.03%	0.02%	0.01%
120	41.12%	0.67%	0.06%	0.06%	0.06%	-0.04%	-0.04%
130	43.88%	0.81%	0.10%	0.10%	0.11%	-0.09%	-0.09%
140	47.18%	0.98%	0.15%	0.16%	0.16%	-0.13%	-0.13%
150	50.81%	1.18%	0.20%	0.21%	0.22%	-0.18%	-0.18%

Table A.8: Approximation errors for the 1 month smile case. Δ P is the difference between the model implied volatility and the exact one. Positive figures indicate overpricing by the model.

Strike	Exact imp vol	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
40	46.81%	-0.42%	-0.14%	-0.14%	-0.14%	-0.06%	-0.02%
50	44.26%	-0.34%	-0.12%	-0.12%	-0.12%	-0.08%	-0.04%
60	41.95%	-0.25%	-0.09%	-0.09%	-0.09%	-0.09%	-0.05%
70	39.86%	-0.18%	-0.07%	-0.07%	-0.07%	-0.09%	-0.05%
80	37.97%	-0.13%	-0.06%	-0.05%	-0.05%	-0.08%	-0.05%
90	36.26%	-0.08%	-0.04%	-0.04%	-0.04%	-0.07%	-0.04%
100	34.72%	-0.04%	-0.03%	-0.03%	-0.03%	-0.06%	-0.03%
110	33.31%	-0.01%	-0.02%	-0.02%	-0.02%	-0.05%	-0.02%
120	32.05%	0.01%	-0.02%	-0.01%	-0.01%	-0.04%	-0.02%
130	30.90%	0.03%	-0.01%	-0.01%	-0.01%	-0.04%	-0.02%
140	29.86%	0.04%	-0.01%	-0.01%	-0.01%	-0.04%	-0.02%
150	28.93%	0.05%	0.00%	0.00%	0.00%	-0.04%	-0.02%
160	28.08%	0.06%	0.00%	0.00%	0.00%	-0.04%	-0.02%
170	27.31%	0.06%	0.00%	0.00%	0.00%	-0.04%	-0.02%
180	26.61%	0.07%	0.00%	0.00%	0.00%	-0.04%	-0.03%
190	25.98%	0.07%	0.00%	0.00%	0.00%	-0.05%	-0.03%
200	25.41%	0.07%	0.00%	0.00%	0.00%	-0.05%	-0.03%

Table A.9: Approximation errors for the 1 year skew case. Δ P is the difference between the model implied volatility and the exact one. Positive figures indicate overpricing by the model.

Strike	Exact imp vol	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
40	38.26%	0.26%	-0.48%	-0.45%	-0.44%	-0.47%	-0.47%
50	33.19%	0.79%	-0.19%	-0.17%	-0.17%	-0.35%	-0.35%
60	29.06%	0.96%	-0.06%	-0.04%	-0.04%	-0.23%	-0.23%
70	26.06%	0.91%	-0.02%	-0.01%	0.00%	-0.13%	-0.13%
80	24.14%	0.81%	-0.03%	-0.01%	-0.01%	-0.05%	-0.05%
90	23.12%	0.75%	-0.02%	-0.01%	-0.01%	0.03%	0.02%
100	22.82%	0.74%	-0.02%	-0.01%	0.00%	0.06%	0.06%
110	23.06%	0.77%	0.00%	0.01%	0.01%	0.05%	0.05%
120	23.71%	0.82%	0.00%	0.01%	0.02%	0.01%	0.00%
130	24.63%	0.88%	0.01%	0.03%	0.03%	-0.03%	-0.04%
140	25.73%	0.95%	0.02%	0.04%	0.04%	-0.07%	-0.07%
150	26.94%	1.00%	0.03%	0.05%	0.05%	-0.10%	-0.10%
160	28.20%	1.04%	0.03%	0.05%	0.05%	-0.13%	-0.13%
170	29.49%	1.03%	0.01%	0.03%	0.03%	-0.16%	-0.17%
180	30.76%	1.01%	-0.02%	0.00%	0.01%	-0.19%	-0.20%
190	32.00%	0.96%	-0.06%	-0.04%	-0.03%	-0.22%	-0.23%
200	33.19%	0.88%	-0.10%	-0.08%	-0.07%	-0.25%	-0.25%

Table A.10: Approximation errors for the 1 year smile case. Δ P is the difference between the model implied volatility and the exact one. Positive figures indicate overpricing by the model.

Strike	Exact imp vol	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
30	28.61%	-0.29%	0.00%	0.00%	0.00%	0.02%	0.02%
40	28.19%	-0.30%	-0.01%	-0.01%	-0.01%	0.00%	0.00%
50	27.79%	-0.30%	-0.02%	-0.01%	-0.01%	-0.02%	-0.02%
60	27.41%	-0.29%	-0.02%	-0.01%	-0.01%	-0.03%	-0.03%
70	27.05%	-0.28%	-0.02%	-0.01%	-0.01%	-0.04%	-0.04%
80	26.70%	-0.27%	-0.02%	-0.01%	-0.01%	-0.04%	-0.04%
90	26.38%	-0.26%	-0.01%	-0.01%	-0.01%	-0.05%	-0.05%
100	26.07%	-0.24%	-0.01%	-0.01%	-0.01%	-0.04%	-0.04%
110	25.77%	-0.23%	-0.01%	0.00%	0.00%	-0.04%	-0.04%
120	25.49%	-0.21%	0.00%	0.00%	0.00%	-0.03%	-0.03%
130	25.22%	-0.19%	0.00%	0.01%	0.01%	-0.03%	-0.03%
140	24.97%	-0.18%	0.01%	0.01%	0.01%	-0.02%	-0.02%
150	24.72%	-0.16%	0.01%	0.01%	0.01%	-0.01%	-0.01%
160	24.49%	-0.15%	0.02%	0.02%	0.02%	-0.01%	-0.01%
170	24.27%	-0.13%	0.02%	0.02%	0.02%	0.00%	0.00%
180	24.07%	-0.12%	0.02%	0.03%	0.03%	0.00%	0.00%
190	23.87%	-0.10%	0.03%	0.03%	0.03%	0.01%	0.01%
200	23.68%	-0.09%	0.03%	0.03%	0.03%	0.01%	0.01%
210	23.50%	-0.08%	0.03%	0.04%	0.04%	0.01%	0.01%
220	23.33%	-0.06%	0.04%	0.04%	0.04%	0.02%	0.02%
230	23.17%	-0.05%	0.04%	0.04%	0.04%	0.02%	0.02%
240	23.01%	-0.04%	0.04%	0.05%	0.05%	0.02%	0.02%
250	22.87%	-0.03%	0.05%	0.05%	0.05%	0.02%	0.02%
260	22.73%	-0.02%	0.05%	0.05%	0.05%	0.02%	0.02%
270	22.59%	-0.01%	0.05%	0.05%	0.05%	0.03%	0.02%
280	22.47%	0.00%	0.05%	0.05%	0.06%	0.03%	0.03%
290	22.35%	0.01%	0.06%	0.06%	0.06%	0.03%	0.03%
300	22.23%	0.02%	0.06%	0.06%	0.06%	0.03%	0.03%

Table A.11: Approximation errors for the 10 year skew case. Δ P is the difference between the model implied volatility and the exact one. Positive figures indicate overpricing by the model.

Strike	Exact imp vol	ΔP1	ΔP2	ΔP3	ΔP4	ΔP5	ΔP6
30	23.41%	-0.13%	-0.47%	-0.45%	-0.44%	-0.26%	-0.27%
40	22.07%	0.23%	-0.35%	-0.33%	-0.32%	-0.21%	-0.21%
50	20.80%	0.56%	-0.23%	-0.21%	-0.20%	-0.12%	-0.13%
60	19.77%	0.79%	-0.14%	-0.12%	-0.11%	-0.04%	-0.05%
70	19.02%	0.93%	-0.09%	-0.06%	-0.05%	0.02%	0.01%
80	18.53%	1.01%	-0.06%	-0.03%	-0.02%	0.07%	0.07%
90	18.28%	1.04%	-0.04%	-0.02%	-0.01%	0.11%	0.10%
100	18.20%	1.05%	-0.04%	-0.01%	0.00%	0.12%	0.11%
110	18.27%	1.05%	-0.04%	-0.01%	0.00%	0.11%	0.11%
120	18.43%	1.03%	-0.04%	-0.02%	-0.01%	0.10%	0.09%
130	18.66%	1.00%	-0.06%	-0.03%	-0.02%	0.07%	0.06%
140	18.93%	0.95%	-0.07%	-0.04%	-0.04%	0.04%	0.04%
150	19.23%	0.90%	-0.09%	-0.06%	-0.05%	0.02%	0.01%
160	19.55%	0.84%	-0.11%	-0.08%	-0.08%	-0.01%	-0.02%
170	19.87%	0.78%	-0.13%	-0.11%	-0.10%	-0.04%	-0.04%
180	20.19%	0.71%	-0.16%	-0.13%	-0.13%	-0.06%	-0.07%
190	20.50%	0.64%	-0.18%	-0.16%	-0.15%	-0.08%	-0.09%
200	20.80%	0.57%	-0.21%	-0.19%	-0.18%	-0.10%	-0.11%
210	21.08%	0.51%	-0.24%	-0.21%	-0.21%	-0.12%	-0.13%
220	21.35%	0.44%	-0.26%	-0.24%	-0.23%	-0.14%	-0.15%
230	21.60%	0.37%	-0.29%	-0.26%	-0.26%	-0.16%	-0.16%
240	21.84%	0.31%	-0.31%	-0.29%	-0.28%	-0.17%	-0.18%
250	22.07%	0.25%	-0.33%	-0.31%	-0.30%	-0.18%	-0.19%
260	22.27%	0.19%	-0.35%	-0.33%	-0.32%	-0.20%	-0.20%
270	22.47%	0.14%	-0.37%	-0.35%	-0.34%	-0.21%	-0.21%
280	22.65%	0.09%	-0.39%	-0.36%	-0.36%	-0.21%	-0.22%
290	22.81%	0.05%	-0.40%	-0.38%	-0.37%	-0.22%	-0.23%
300	22.97%	0.00%	-0.42%	-0.39%	-0.38%	-0.23%	-0.23%

Table A.12: Approximation errors for the 10 year smile case. Δ P is the difference between the model implied volatility and the exact one. Positive figures indicate overpricing by the model.

Bibliography

ANDERSEN, LEIF, & ANDREASEN, JESPER. 2000. Volatility skews and extensions of the LIBOR market model. *Applied Mathematical Finance*, **7**, 1–32.

ANDERSEN, LEIF, & BROTHERTON-RATCLIFFE, RUPERT. 1998. The equity option volatility smile: An implicit finite-difference approach. *Journal of Computational Finance*, **1**, 5–38.

ANDERSEN, LEIF B.G., & BROTHERTON-RATCLIFFE, RUPERT. 2001. Extended LIBOR market models with stochastic volatility. Working paper. Gen Re Securities.

APPLEBAUM, DAVID. 2004. *Lévy processes and stochastic calculus.* Cambridge University Press.

AVELLANEDA, MARCO, FRIEDMAN, CRAIG, HOLMES, RICHARD, & SAMPERI, DOMINICK. 1997. Calibrating volatility surfaces via relative entropy minimization. *Applied Mathematical Finance*, **4**(1), 37–64.

BAKSHI, G., & MADAN, DILIP B. 2000. Spanning and derivative security valuation. *Journal of Financical Economics*, **55**, 205–238.

BALLAND, PHILIPPE. 2002. Deterministic implied volatility models. *Quantitative Finance*, **2**, 31–44.

BARLE, STANKO, & CAKICI, NUSRET. 1995. Growing a smiling tree. *RISK journal*, **10**(8), 76–81.

BARNDORFF-NIELSEN, OLE E. 1998. Processes of normal inverse Gaussian type. *Finance and Stochastics*, **2**(1), 41–68.

BARNDORFF-NIELSEN, OLE E., NICOLATO, ELISA, & SHEPHARD, NEIL. 2002. Some recent developments in stochastic volatility modelling. *Quantitative Finance*, **2**, 11–23.

BATES, DAVID S. 1996. Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies*, **9**, 69–107.

BERESTYCKI, HENRI, BUSCA, JÉRÔME, & FLORENT, IGOR. 2002. Asymptotics and calibration of local volatility models. *Quantitative Finance*, **2**(1), 61–69.

BERTOIN, JEAN. 1998. *Lévy processes.* Cambridge University Press.

BINGHAM, NICK H., & KIESEL, RÜDIGER. 2004. *Risk-neutral valuation.* Springer-Verlag, Berlin, Heidelberg, New York.

BJÖRK, TOMAS, DI MASI, GIOVANNI, KABANOV, YURI, & RUNGGALDIER, WOLFGANG. 1997. Towards a general theory of bond markets. *Finance and Stochastics*, **1**, 141–174.

BLACK, FISCHER. 1976. The pricing of commodity contracts. *Journal of Financial Economics*, **3**, 167–179.

BLACK, FISCHER, & SCHOLES, MYRON S. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy*, **81**, 637–659.

BOUCHOUEV, ILIA, & ISAKOV, VICTOR. 1997. The inverse problem of option pricing. *Inverse Problems*, **13**, 69–107.

BOUCHOUEV, ILIA, & ISAKOV, VICTOR. 1999. Uniqueness, stability and numerical methods for the inverse problem that arises in financial markets. *Inverse Problems*, **15**, 95–116.

BRACE, ALAN, GATAREK, DARIUSZ, & MUSIELA, MAREK. 1996. The market model of interest rate dynamics. *Mathematical Finance*, **7**, 127–154.

BRACE, ALAN, MUSIELA, MAREK, & SCHLÖGL, LUTZ. 1998. A simulation algorithm based on measure relationships in the lognormal market models. Working paper. UNSW.

BRACE, ALAN, GOLDYS, BENJAMIN, KLEBANER, FIMA, & WOMERSLEY, ROBERT. 2001. Market model of implied volatility with application to BGM. Working paper. UNSW.

BRACE, ALAN, GOLDYS, BENJAMIN, VAN DER HOEK, JOHN, & WOMERSLEY, ROBERT. 2002. Markovian models in the stochastic implied volatility framework. Working paper. UNSW.

BREEDEN, DOUGLAS T., & LITZENBERGER, ROBERT H. 1978. Prices of contingent claims implied in option prices. *Journal of Business*, **51**, 621–651.

BRIGO, DAMIANO, & MERCURIO, FABIO. 2001. *Interest rate models - theory and practice.* Springer-Verlag, Berlin, Heidelberg, New York.

BRIGO, DAMIANO, & MERCURIO, FABIO. 2003. Analytical pricing of the smile in a forward LIBOR market model. *Quantitative Finance*, **3**(1), 15–27.

BRIGO, DAMIANO, MERCURIO, FABIO, & SARTORELLI, GIULIO. 2003. Alternative asset price dynamics and volatility smile. *Quantitative Finance*, **3**(3), 173–183.

CARR, PETER, & WU, LIUREN. 2003. Finite moment log stable process and option pricing. *Journal of Finance*, **58**(2), 753–777.

CARR, PETER, GEMAN, HÉLYETTE, MADAN, DILIP B., & YOR, MARC. 2002. The fine structure of asset returns: An empirical investigation. *Journal of Business*, **75**(2), 305–332.

CARR, PETER, GEMAN, HÉLYETTE, MADAN, DILIP B., & YOR, MARC. 2003. Stochastic volatility for Lévy processes. *Mathematical Finance*, **13**(3), 345–382.

CHAN, TERENCE. 1999. Pricing contingent claims on stocks driven by Lévy processes. *Annals of Applied Probability*, **9**(2), 504–528.

CONT, RAMA, & DA FONSECA, JOSÉ. 2002. Dynamics of implied volatility surfaces. *Quantitative Finance*, **2**(1), 45–60.

CONT, RAMA, & TANKOV, PETER. 2004. *Financial modelling with jump processes.* Chapman & Hall/CRC.

DAGLISH, TOBY, HULL, JOHN C., & SUO, WULIN. 2003. Volatility surfaces: Theory, rules of thumb, and empirical evidence. Working paper. University of Toronto and Queen's University.

DEMPSTER, MICHAEL A.H., & RICHARDS, DARREN G. 2000. Pricing American options fitting the smile. *Mathematical Finance*, **10**(2), 157–177.

DERMAN, EMANUEL. 1999. Regimes of volatility. Goldman Sachs Quantitative Strategies Research Notes.

DERMAN, EMANUEL. 2003. Laughter in the dark – The problem of the volatility smile. Working paper.

DERMAN, EMANUEL, & KANI, IRAJ. 1994. Riding on a smile. *RISK journal*, **7**(2), 32–39.

DERMAN, EMANUEL, & KANI, IRAJ. 1998. Stochastic implied trees: Arbitrage pricing with stochastic term- and strike structure of volatility. *International Journal of Theoretical and Applied Finance*, **1**, 61–110.

DERMAN, EMANUEL, KANI, IRAJ, & ZOU, JOSEPH Z. 1995. The local volatility surface - Unlocking the information in index option prices. Goldman Sachs Quantitative Strategies Research Notes.

DUMAS, BERNARD, FLEMING, JEFFERSON, & WHALEY, ROBERT E. 1998. Implied volatility functions - Empirical tests. *Journal of Finance*, **53**(6), 2059–2106.

DUPIRE, BRUNO. 1993. Pricing and hedging with smiles. Working paper. Paribas Capital Markets.

DUPIRE, BRUNO. 1994. Pricing with a smile. *RISK journal*, 18–20.

DUPIRE, BRUNO. 1997. Pricing and hedging with smiles. *Pages 103–111 of:* DEMPSTER, MICHAEL A.H., & PLISKA, STANLEY R. (eds), *Mathematics of derivative securities.* Publications of the Newton Institute. Cambridge University Press.

EBERLEIN, ERNST, & KELLER, ULRICH. 1995. Hyperbolic distributions in finance. *Bernoulli*, **1**, 281–299.

EBERLEIN, ERNST, & ÖZKAN, FEHMI. 2004. Lévy processes in credit risk and market models. Working paper. Albert-Ludwigs-Universität Freiburg i.Br.

ESSER, ANGELIKA, & SCHLAG, CHRISTIAN. 2001. A note on forward and backward partial differential equations for derivative contracts with forwards as underlyings. Working paper. Goethe University Frankfurt am Main.

GATHERAL, JIM. 2000. Rational shapes of the volatility surface. Presentation at the RISK conference.

GATHERAL, JIM. 2003. Case studies in financial modelling course notes. Courant Institute of Mathematical Sciences. Fall Term, 2003.

GEMAN, HÉLYETTE. 2002. Pure jump Lévy processes for asset price modelling. *Journal of Banking and Finance*, **26**, 1297–1316.

GEMAN, HÉLYETTE, EL KAROUI, NICOLE, & ROCHET, JEAN-CHARLES. 1995. Changes of numéraire, changes of probability measure and option pricing. *Journal of Applied Probability*, 443–458.

GEMAN, HÉLYETTE, MADAN, DILIP B., & YOR, MARC. 2001. Time changes for Lévy processes. *Mathematical Finance*, **11**(1), 79–96.

GLASSERMAN, PAUL, & KOU, STEVEN. 2003. The term structure of simple forward rates with jump risk. *Mathematical Finance*, July, 383–410.

GLASSERMAN, PAUL, & MERENER, NICOLAS. 2003a. Cap and swaption approximations in LIBOR market models with jumps. *Journal of Computational Finance*, **7**(1), 1–36.

GLASSERMAN, PAUL, & MERENER, NICOLAS. 2003b. Numerical solution of jump-diffusion LIBOR market models. *Finance and Stochastics*, **7**(1), 1–27.

GLASSERMAN, PAUL, & ZHAO, XIAOLIANG. 1999. Fast greeks by simulation in forward LIBOR models. *Journal of Computational Finance*, **3**(1), 5–39.

GLASSERMAN, PAUL, & ZHAO, XIAOLIANG. 2000. Arbitrage-free discretization of lognormal forward LIBOR and swap rate models. *Finance and Stochastics*, **4**, 35–68.

HAGAN, PATRICK S., & WOODWARD, DIANA E. 1999. Equivalent Black volatilities. *Applied Mathematical Finance*, **6**, 147–157.

HAGAN, PATRICK S., KUMAR, DEEP, LESNIEWSKI, ANDREW S., & WOODWARD, DIANA E. 2002. Managing smile risk. *Wilmott magazine*, Sept., 84–108.

HESTON, STEVEN. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, **6**(2), 327–343.

HULL, JOHN C., & WHITE, ALAN. 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance*, **42**(6), 281–300.

HULL, JOHN C., & WHITE, ALAN. 2000. Forward rate volatilities, swap rate volatilities, and implementation of the LIBOR market model. *Journal of Fixed Income*, **10**(3), 46–62.

HUNT, PHIL J., & KENNEDY, JOANNE E. 2000. *Financial derivatives in theory and practice.* Wiley & Sons Ltd.

JACKSON, NICOLAS, SÜLI, ENDRE, & HOWISON, SAM. 1999. Computation of deterministic volatility surfaces. *Journal of Computational Finance*, **2**(2), 5–32.

JACOD, JEAN, & SHIRYAEV, ALBERT NIKOLAEVICH. 1987. *Limit theorems for stochastic processes.* Springer-Verlag, Berlin, Heidelberg, New York.

JAMSHIDIAN, FARSHID. 1997. LIBOR and swap market models and measures. *Finance and Stochastics*, **1**, 261–291.

JARROW, ROBERT, LI, HAITAO, & ZHAO, FENG. 2003. Interest rate caps 'smile' too! But can the LIBOR market models capture it? Working paper. Cornell University.

JIANG, LISHANG, CHEN, QIHONG, WANG, LIJUN, & ZHANG, JIN E. 2003. A new well-posed algorithm to recover implied local volatility. *Quantitative Finance*, **3**, 451–457.

JOSHI, MARK S., & REBONATO, RICCARDO. 2001. A stochastic-volatility, displaced-diffusion extension of the LIBOR market model. Working paper. QUARC.

JOSHI, MARK S., & REBONATO, RICCARDO. 2003. A displaced-diffusion stochastic volatility LIBOR market model: Motivation, definition and implementation. *Quantitative Finance*, **3**(6), 485–469.

KLEBANER, FIMA. 2002. Option price when the stock is a semimartinagle. *Electronic Communications in Probability*, **7**, 79–83.

KLEBANER, FIMA. 2003. Correction to: Option price when the stock is a semimartingale. *Electronic Communications in Probability*, **8**, C.1.

KOU, STEVEN G. 2002. A jump-diffusion model for option pricing. *Management Science*, **48**(8), 1086–1101.

LAGNADO, RONALD, & OSHER, STANLEY. 1997. A technique for calibrating derivative security pricing models: Numerical solution of an inverse problem. *Journal of Computational Finance*, **1**(1), 13–25.

LEDOIT, OLIVIER, SANTA-CLARA, PEDRO, & YAN, SHU. 2002. Relative pricing of options with stochastic volatility. Working paper.

LEE, ROGER. 2002. Implied volatility: Statics, dynamics, and probabilistic interpretation. Working paper. Stanford University.

LI, YANMIN. 2001. A new algorithm for constructing implied binomial trees: Does the implied model fit any volatility smile? *Journal of Computational Finance*, **4**(2), 69–95.

MADAN, DILIP B., CARR, PETER P., & CHANG, ERIC C. 1998. The Variance Gamma process and option pricing. *European Finance Review*, **2**, 79–105.

MERTON, ROBERT. 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science*, **4**(1), 141–183.

MERTON, ROBERT. 1976. Option pricing with discontinuous returns. *Journal of Financial Economics*, **3**, 145–166.

MILTERSEN, KRISTIAN R., SANDMANN, KLAUS, & SONDERMANN, DIETER. 1997. Closed form solutions for term-structure derivatives with log-normal interest rates. *Journal of Finance*, **52**, 409–430.

ØKSENDAL, BERNT. 2000. *Stochastic differential equations - An introduction with applications.* 5th edn. Springer-Verlag, Berlin, Heidelberg, New York.

ÖZKAN, FEHMI. 2002. *Lévy processes in credit risk and market models.* Ph.D. thesis, Albert-Ludwigs-Universität Freiburg i.Br.

PELSSER, ANTOON. 2000. *Efficient models for valuing interest rate derivatives.* Springer-Verlag, Berlin, Heidelberg, New York.

PRAUSE, KARSTEN. 1999. *The generalized hyperbolic model: Estimation, financial derivatives, and risk measures.* Ph.D. thesis, Albert-Ludwigs-Universität Freiburg i.Br.

PROTTER, PHILIP E. 2003. *Stochastic integration and differential equations.* Springer-Verlag, Berlin, Heidelberg, New York.

RAIBLE, SEBASTIAN. 2000. *Lévy processes in finance: Theory, numerics and empirical facts.* Ph.D. thesis, Albert-Ludwigs-Universität Freiburg i.Br.

REBONATO, RICCARDO. 2002. *Modern pricing of interest-rate derivatives.* Princeton University Press, Princeton and Oxford.

REBONATO, RICCARDO. 2003. Term structure models: A review. Working paper. University of Oxford.

REBONATO, RICCARDO. 2004. *Volatility and correlation: The perfect hedger and the fox.* Second edn. John Wiley & Sons Ltd.

RUBINSTEIN, MARK. 1994. Implied binomial trees. *Journal of Finance*, **49**(3), 771–818.

SATO, KEN-ITI. 1999. *Lévy processes and infinitely divisible distributions.* Cambridge University Press.

SCHOUTENS, WIM. 2003. *Lévy processes in finance – Pricing financial derivatives.* John Wiley & Sons Ltd.

SCHOUTENS, WIM, SIMONS, ERWIN, & TISTAERT, JURGEN. 2003. A perfect calibration! Now what? Working paper.

SHREVE, STEVEN E. 2004. *Stochastic calculus for finance II: Continuous-time models.* Springer-Verlag, Berlin, Heidelberg, New York.

ZÜHLSDORFF, CHRISTIAN. 2002. Extended market models with affine and quadratic volatility. Discussion paper. University of Bonn.