

Editorial Board

M. Avellaneda

G. Barone-Adesi

M. Broadie

M.H.A. Davis

E. Derman

C. Klüppelberg

E. Kopp

W. Schachermayer

Springer Finance

Springer Finance is a programme of books aimed at students, academics and practitioners working on increasingly technical approaches to the analysis of financial markets. It aims to cover a variety of topics, not only mathematical finance but foreign exchanges, term structure, risk management, portfolio theory, equity derivatives, and financial economics.

- Ammann M.*, Credit Risk Valuation: Methods, Models, and Application (2001)
Back K., A Course in Derivative Securities: Introduction to Theory and Computation (2005)
Barucci E., Financial Markets Theory. Equilibrium, Efficiency and Information (2003)
Bielecki T.R. and Rutkowski M., Credit Risk: Modeling, Valuation and Hedging (2002)
Bingham N.H. and Kiesel R., Risk-Neutral Valuation: Pricing and Hedging of Financial Derivatives (1998, 2nd ed. 2004)
Brigo D. and Mercurio F., Interest Rate Models: Theory and Practice (2001)
Buff R., Uncertain Volatility Models-Theory and Application (2002)
Dana R.A. and Jeanblanc M., Financial Markets in Continuous Time (2002)
Deboeck G. and Kohonen T. (Editors), Visual Explorations in Finance with Self-Organizing Maps (1998)
Elliott R.J. and Kopp P.E., Mathematics of Financial Markets (1999, 2nd ed. 2005)
Fengler M., Semiparametric Modeling of Implied Volatility (2005)
Geman H., Madan D., Pliska S.R. and Vorst T. (Editors), Mathematical Finance-Bachelier Congress 2000 (2001)
Gundlach M., Lehrbass F. (Editors), CreditRisk⁺ in the Banking Industry (2004)
Kellerhals B.P., Asset Pricing (2004)
Külpmann M., Irrational Exuberance Reconsidered (2004)
Kwok Y.-K., Mathematical Models of Financial Derivatives (1998)
Malliavin P. and Thalmaier A., Stochastic Calculus of Variations in Mathematical Finance (2005)
Meucci A., Risk and Asset Allocation (2005)
Pelsser A., Efficient Methods for Valuing Interest Rate Derivatives (2000)
Prigent J.-L., Weak Convergence of Financial Markets (2003)
Schmid B., Credit Risk Pricing Models (2004)
Shreve S.E., Stochastic Calculus for Finance I (2004)
Shreve S.E., Stochastic Calculus for Finance II (2004)
Yor, M., Exponential Functionals of Brownian Motion and Related Processes (2001)
Zagst R., Interest-Rate Management (2002)
Ziegler A., Incomplete Information and Heterogeneous Beliefs in Continuous-time Finance (2003)
Ziegler A., A Game Theory Analysis of Options (2004)
Zhu Y.-L., Wu X., Chern I.-L., Derivative Securities and Difference Methods (2004)

Matthias R. Fengler

Semiparametric Modeling of Implied Volatility

 Springer

Matthias R. Fengler
Equity Derivatives Group
Sal. Oppenheim jr. & Cie.
Untermainanlage 1
60329 Frankfurt
Germany
E-mail: matthias.fengler@oppenheim.de

Mathematics Subject Classification (2000): 62G08, 62G05, 62H25

JEL classification: G12, G13

This book is based on the author's dissertation accepted on 28 June 2004 at the Humboldt-Universität zu Berlin.

Library of Congress Control Number: 2005930475

ISBN-10 3-540-26234-2 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-26234-3 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable for prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springeronline.com

© Springer-Verlag Berlin Heidelberg 2005

Printed in The Netherlands

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: by the author and TechBooks using a Springer L^AT_EX macro package

Cover design: *design & production*, Heidelberg

Printed on acid-free paper SPIN: 11496786 41/TechBooks 5 4 3 2 1 0

LE MONDE INSTABLE

Le monde en vne isle porté
Sur la mer tant esmeue et rogue,
Sans seur gouuernal nage et vogue,
Monstrant son instabilité.

Corrozet (1543)

quoted from Henkel and Schöne (1996)

Acknowledgements

This book has benefitted a lot from suggestions and comments of colleagues, fellow students and friends whom I wish to thank at this place. At first rate, I thank Wolfgang Härdle. He directed my interest to implied volatilities and made me familiar with non- and semiparametric modeling in Finance. Without him, his encouragement and advise this work would not exist. Furthermore, I like to thank Vladimir Spokoiny, in particular for his comments during my talks in the Seminar for Mathematical Statistics at the WIAS, Berlin.

This work is in close context with essays I have written with a number of coauthors. Above all, I thank Enno Mammen: the cooperation in semiparametric modeling has been highly instructive and fruitful for me. In this regard, I also thank Qihua Wang.

For an unknown number of helpful discussions or proofreading my thanks go to Peter Bank, Michal Benko, Szymon Borak, Kai Detlefsen, Erhard and Martin Fengler, Patrick Herbst, Zdeněk Hlávka, Torsten Kleinow, Danilo Mercurio and Marlene Müller and to all contemporary and former members of the ISE and CASE for the inspiring working environment they generated there.

Finally, I wish to thank the members of my family non explicitly mentioned up to now, Stephanus and especially my mother Brigitte Fengler and Georgia Mavrodi who in their ways did all their best to support me and the project at its different stages.

I gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft in having been a member of the Sonderforschungsbereich 373 *Quantifikation und Simulation ökonomischer Prozesse* at the Humboldt-Universität zu Berlin.

Berlin, May 2005

Matthias R. Fengler

Frequently Used Notation

Abbreviation or symbol	Explanation
ATM	at-the-money
BS	Black and Scholes (1973)
cdf	cumulative distribution function
C_t	price of a call option at time t
C_t^{BS}	Black-Scholes price of a call option at time t
$\mathcal{C}(\mathcal{A})$	the continuous functions $f : \mathcal{A} \rightarrow \mathbb{R}$
$\mathcal{C}^k(\mathcal{A})$	functions in $\mathcal{C}(\mathcal{A})$ with continuous derivatives up to order k
$\mathcal{C}^{k,l}(\mathbb{R} \times \mathbb{R})$	the functions $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ which are \mathcal{C}^k w.r.t. the first and \mathcal{C}^l w.r.t. the second argument
$\text{Cov}(X, Y)$	covariance of two random variables X and Y
CPC(A)	common principal component (analysis)
δ	dividend yield
δ_{x_0}	Dirac delta function defined by the property: $\int f(x) \delta_{x_0}(x) dx = f(x_0)$ for a smooth function f
$E(X)$	expected value of the random variable X
F_t	forward or futures price of an asset at time t
\mathcal{F}_t	filtration, the information set generated by the information available up to time t
\mathbf{I}_p	$p \times p$ unity matrix
IV	implied volatility
IVS	implied volatility surface
ITM	in-the-money
$\mathbf{1}(\mathcal{A})$	indicator function of the set \mathcal{A}
K	exercise price

$K(\cdot)$	kernel function: continuous, bounded and symmetric real function satisfying $\int K(u) du = 1$
κ_f	forward or futures moneyness: $\kappa_f \stackrel{\text{def}}{=} K/F_t$
LVS	local volatility surface
μ	mean of a random variable
$N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	normal distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$
OTM	out-of-the-money
\mathcal{O}	$\alpha_n = \mathcal{O}(\beta_n)$ means: $\lim_{n \rightarrow \infty} \frac{\alpha_n}{\beta_n} \rightarrow 0$
\mathcal{O}	$\alpha_n = \mathcal{O}(\beta_n)$ means: $\lim_{n \rightarrow \infty} \frac{\alpha_n}{\beta_n} \rightarrow \text{some constant}$
pdf	probability density function
pCPC(q)	partial CPC model of order q
P_t	price of a put option at time t
PCA	principal component analysis
$P(\mathcal{A})$	probability of the set \mathcal{A} , objective measure
PDE	partial differential equation
Q	a risk neutral measure
r	interest rate
\mathbb{R}^d	d -dimensional Euclidian space, $\mathbb{R} = \mathbb{R}^1$
\mathbb{R}^+	the non-negative real numbers
S_t	price of a stock at time t
SDE	stochastic differential equation
$\boldsymbol{\Sigma}$	covariance matrix
t	time
T	expiry date of a financial contract
τ	$\tau \stackrel{\text{def}}{=} T - t$, time to maturity of an option or a forward
$\text{Var}(X)$	variance of the random variable X
W_t	Brownian motion at time t
\overline{W}_t	Brownian motion under the risk neutral measure at time t
$\varphi(x)$	pdf of the normal distribution: $\varphi(x) \stackrel{\text{def}}{=} \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$
$\Phi(u)$	cdf of a normal random variable: $\Phi(u) \stackrel{\text{def}}{=} \int_{-\infty}^u \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx$
$\stackrel{\text{def}}{=}$	is defined as
\sim	if $X \sim D$, the random variable X has the distribution D
$\xrightarrow{\mathcal{L}}$	converges in distribution to
\xrightarrow{p}	converges in probability to
$(X)^+$	$(X)^+ \stackrel{\text{def}}{=} \max(X, 0)$
$\langle X \rangle_t$	quadratic variation process of the stochastic process X
$\langle X, Y \rangle_t$	covariation process of the stochastic processes X and Y
$ x $	absolute value of the scalar x
$ \mathbf{X} $	determinant of the matrix \mathbf{X}
\mathbf{X}^\top	transpose of the matrix \mathbf{X}
$\text{tr } \mathbf{X}$	trace of the matrix \mathbf{X}
$\langle f, g \rangle$	inner product of the functions f and g

In this book, we will mainly employ three concepts of volatility based on the following stochastic differential equation for the asset price process:

$$\frac{dS_t}{S_t} = \mu(S_t, t) dt + \sigma(S_t, t, \cdot) dW_t .$$

These concepts are in particular:

Instantaneous	Implied	Local
—— volatility ——		
$\sigma(S_t, t, \cdot)$	$\hat{\sigma}_t(K, T)$	$\sigma_{K,T}(S_t, t)$

Instantaneous volatility measures the instantaneous standard deviation of the return process of the log-asset price. It depends on the current level of the asset price S_t , time t and possibly on other state variables abbreviated with ‘ \cdot ’.

Implied volatility is the BS option price implied measure of volatility. It is the volatility parameter that equates the BS price and a particular observed market price of an option. Thus, it depends on the strike K , the expiry date T and time t .

Local volatility is the expected instantaneous volatility conditional on a particular level of the asset price $S_T = K$ at $t = T$. If the instantaneous volatility is a deterministic function in S_t and t , i.e. can be written as $\sigma(S_t, t)$, then $\sigma_{K,T}(S_t, t) = \sigma(K, T)$.

The term **volatility** is reserved for objects of the kind σ and $\hat{\sigma}$, while their squared counterparts σ^2 and $\hat{\sigma}^2$ are called **variance**.

Contents

1	Introduction	1
2	The Implied Volatility Surface	9
2.1	The Black-Scholes Model	9
2.2	The Self-Financing Replication Strategy	11
2.3	Risk Neutral Pricing	12
2.4	The BS Formula and the Greeks	14
2.5	The IV Smile	19
2.6	Static Properties of the Smile Function	27
2.6.1	Bounds on the Slope	27
2.6.2	Large and Small Strike Behavior	28
2.7	General Regularities of the IVS	30
2.7.1	Static Stylized Facts	30
2.7.2	DAX Index IV between 1995 and 2001	33
2.8	Relaxing the Constant Volatility Case	34
2.8.1	Deterministic Volatility	35
2.8.2	Stochastic Volatility	36
2.9	Challenges Arising from the Smile	40
2.9.1	Hedging and Risk Management	40
2.9.2	Pricing	42
2.10	IV as Predictor of Realized Volatility	42
2.11	Why Do We Smile?	43
2.12	Summary	46
3	Smile Consistent Volatility Models	47
3.1	Introduction	47
3.2	The Theory of Local Volatility	49
3.3	Backing the LVS Out of Observed Option Prices	51
3.4	The <i>dual</i> PDE Approach to Local Volatility	54
3.5	From the IVS to the LVS	55
3.6	Asymptotic Relations Between Implied and Local Volatility ...	60

3.7	The Two-Times-IV-Slope Rule for Local Volatility	62
3.8	The K -Strike and T -Maturity Forward Risk-Adjusted Measure	64
3.9	Model-Free (Implied) Volatility Forecasts	66
3.10	Local Volatility Models	67
	3.10.1 Deterministic Implied Trees	67
	3.10.2 Stochastic Implied Trees	80
	3.10.3 Reconstructing the LVS	84
3.11	Excellent Fit, but...: the Delta Problem	88
3.12	Stochastic IV Models	91
3.13	Summary	94
4	Smoothing Techniques	97
4.1	Introduction	97
4.2	Nadaraya-Watson Smoothing	99
	4.2.1 Kernel Functions	99
	4.2.2 The Nadaraya-Watson Estimator	100
4.3	Local Polynomial Smoothing	102
4.4	Bandwidth Selection	104
	4.4.1 Theoretical Framework	104
	4.4.2 Bandwidth Choice in Practice	106
4.5	Least Squares Kernel Smoothing	115
	4.5.1 The LSK Estimator of the IVS	115
	4.5.2 Application of the LSK Estimator	117
4.6	Summary	123
5	Dimension-Reduced Modeling	125
5.1	Introduction	125
5.2	Common Principal Component Analysis	128
	5.2.1 The Family of CPC Models	128
	5.2.2 Estimating Common Eigenstructures	131
	5.2.3 Stability Tests for Eigenvalues and Eigenvectors	134
	5.2.4 CPC Model Selection	138
	5.2.5 Empirical Results	139
5.3	Functional Data Analysis	155
	5.3.1 Basic Set-Up of FPCA	156
	5.3.2 Computing FPCs	157
5.4	Semiparametric Factor Models	160
	5.4.1 The Model	162
	5.4.2 Norming of the Estimates	166
	5.4.3 Choice of Model Parameters	167
	5.4.4 Empirical Analysis	171
	5.4.5 Assessing Prediction Performance	182
5.5	Summary	184

6 Conclusion and Outlook 187

A Description and Preparation of the IV Data 189

 A.1 Preliminaries 189

 A.2 Data Correction Scheme 190

B Some Results from Stochastic Calculus 195

C Proofs of the Results on the LSK IV Estimator 201

 C.1 Proof of Consistency 201

 C.2 Proof of Asymptotic Normality 203

References 207

Index 221

Introduction

Yet that weakness is also its greatest strength. People like the model because they can easily understand its assumptions. The model is often good as a first approximation, and if you can see the holes in the assumptions you can use the model in more sophisticated ways.

Black (1992)

Expected volatility as a measure of risk involved in economic decision making is a key ingredient in modern financial theory: the rational, risk-averse investor will seek to balance the tradeoff between the risk he bears and the return he expects. The more volatile the asset is, i.e. the more it is prone to excessive price fluctuations, the higher will be the expected premium he demands. Markowitz (1959), followed by Sharpe (1964) and Lintner (1965), were among the first to quantify the idea of the simple equation ‘more risk means higher return’ in terms of equilibrium models. Since then, the analysis of volatility and price fluctuations has sparked a vast literature in theoretical and quantitative finance that refines and extends these early models. As the most recent climax of this story, one may see the Nobel prize in Economics granted to Robert Engle in 2003 for his path-breaking work on modeling time-dependent volatility.

Long before this, a decisive turn in the research of volatility was rendered possible with the seminal publication by Black and Scholes (1973) on the pricing of options and corporate liabilities. Their fundamental result, the celebrated Black-Scholes (BS) formula, offers a framework for the valuation of European style derivatives within a simple set of assumptions. Six parameters enter the pricing formula: the current underlying asset price, the strike price, the expiry date of the option, the riskless interest rate, the dividend yield, and a constant volatility parameter that describes the instantaneous standard deviation of the returns of the log-asset price. The application of the formula, however, faces an obstacle: only its first five parameters are known quantities. The last one, the volatility parameter, is not.

An obvious way to respond to this dilemma is to resort to well-established statistical tools and to estimate the volatility parameter from the time series data of the underlying asset. However, there is also a second perspective that the markets and the literature quickly adopted: instead of estimating the volatility for finding an option price, one aims at recovering that volatility

IVS Ticks 20000502




Fig. 1.1. DAX option IVs on 20000502. IV observations are displayed as *black dots*. *Lower left axis* is moneyness and *lower right* time to maturity measured in years

which the market has priced into a given option price observation. To put it in other words, the question is:

what volatility is implied in observed option prices, if the BS model is a valid description of market conditions?

This *reverse* perspective constitutes the concept of the *BS implied volatility*.

A typical picture of implied volatility (IV), as observed on 2nd May, 2000, or 20000502 (a date notation we will adopt from now on) is presented in Fig. 1.1. IV is displayed across different strike prices and expiry dates. Strikes are rescaled in a moneyness metric, where strikes near the current asset price are mapped into the neighborhood of one, and the expiry dates are converted into the time to maturity of the option expressed in years. As is visible, IV exhibits a pronounced curvature across strikes and is also curved across time to maturity, albeit not so much. For a given time to maturity, this function has been named *smile*, and the entire ensemble is called the *implied volatility surface* (IVS). The striking conclusion from a picture like Fig. 1.1 is the clear contradiction to an assumption fundamental to the BS model: instead of being constant, IV is nonlinear in strikes and time to maturity, and – if seen in a sequel of points in time – also time-dependent.

This evident antagonism has been a fruitful starting-point for variations and extensions of this basic pricing model in any direction. At the same time, it does not appear to harm the model itself or the popularity of IV. Nowadays, IV is ubiquitous: it serves as a convenient way of quoting options among

market participants, volatility trading is common practice on trading floors, market models incorporate the risk from fluctuating IVs for hedges, and risk management tools, which are approved by banking regulators to steer the allocation of economic capital, include models of the IVS.

A number of reasons may be put forward for explaining the unrivalled popularity of IV. One of them – already anticipated by the initial words by Fisher Black – can be seen in the set of easy-to-communicate assumptions associated with the BS model. Another, more fundamental reason is that a volatility concept implied from option prices enjoys a particular – if not pivotal – property: as options are bets on the future development of the underlying asset, the key advantage of this *option implied volatility* is the fact that it is a forward looking variable by nature. Thus, unlike volatility measures based on historical data, it should reflect *market expectations* on volatility over the remaining life time of the option. Consequently, the information content of IV and its capability of being a predictor for future asset price volatility has been of primary concern in the literature on IV from the early studies up to now.

Yet, it was only in the recent decade that the finance community recognized that the IVS – aside from being a potential predictor or well-known artefact and curiosity – bears valuable information on the asset price process and its dynamics, and that this information can be exploited in models for the pricing and the hedging of other complex derivatives or positions. This development goes in line with the advent of highly liquid option and futures markets that were established all around the world beginning from the nineteen-nineties. Before this, model calibration and pricing typically relied on historically sampled time series data. This bears the disadvantage that the results are predominantly determined by the price history and that the adjustment to new information is too slow. Unlike time series data, the cross-sectional dimension of option prices across different strikes over a range of time to maturities offers the unique opportunity to directly exploit *instantaneous* data for model calibration.

This breakthrough, initiated by the work of Derman and Kani (1994a), Dupire (1994) and Rubinstein (1994), triggered the literature on *smile consistent pricing*. It led, for instance, to the development of static option replication as a means of hedging or to implied trees as a pricing tool. The challenge for this new approach is that IV cannot be directly used as an input factor, since – as shall be seen in the course of this book – IV is a *global* measure of volatility. Pricing requires a *local* measure of volatility. Hence, at the heart of this theory there is another volatility concept, called *local volatility*. Local volatility, unfortunately, cannot be observed and needs to be extracted from market data, either from option prices or from the IVS. Other modeling approaches formulate IV as an additional stochastic process, that – together with the asset price process – enters the pricing equation of derivatives.

These developments explain why the new focus actuated the interest in refined modeling techniques of the IVS and in the structural analysis of its dynamics. In modeling the IVS, one faces two principal challenges: as is visible

from Fig. 1.1, the estimators are required to provide sufficient *functional flexibility* in order to optimally fit the shape of the IVS. Otherwise, a model bias will ensue. Second, given the high-dimensional complexity of the IVS, *low-dimensional representations* are desirable from a dynamic standpoint. Not only does a low-dimensional representation of the IVS facilitate the practical implementation of any (dynamic) model, it additionally uncovers the structural basis of the data. This will ultimately lead to a better understanding of the IVS as a financial variable. Natural candidates of techniques that meet these key requirements are *non- and semiparametric methods*: they allow for high functional flexibility and parsimonious modeling. Therefore, results from this line of research are of immediate importance when local volatility or stochastic IV models are to be implemented in practice.

The aim of this book is twofold: the first object is to give a thorough treatment of the financial theory on implied and local volatility and smile consistent modeling. Particular attention is given to highlight the cross-relationships between the volatility concepts as shown in Fig. 1.1. The second object is to familiarize the reader with refined non- and semiparametric estimation strategies and dimension reduction methods for functional surfaces and to demonstrate their effectiveness in the field of IV modeling. The majority of results and techniques we discuss are currently available in preprints or published papers, only. In having their applicability in mind, we take care to illustrate them with empirical investigations that underline their use in practice. We believe that in combining the two fields of research – smile consistent modeling and non- and semiparametric estimation techniques – in this way, we can fill a gap among the textbooks at today's disposal.

Writing a book in the mid of two fields of research requires concessions to the breadth each topic can be treated with. Since our emphasis is on financial modeling aspects, we introduce both financial and statistical theory to the extent we deem necessary for the reader to fully appreciate the core concepts of the book. At the same time, we try to keep the book as self-contained as possible in providing an appendix that collects main results from stochastic calculus and statistics. Therefore, general asset pricing theory is introduced only in its basics. For a broader and more general overview on asset pricing theory the reader is referred to classical textbooks such as Björk (1998), Duffie (2001), Föllmer and Schied (2002), Hull (2002), Joshi (2003), or Lipton (2001) to name but a few. The same philosophy applies to the non- and semiparametric methods. Standard books the reader may like to consult in this direction are provided, e.g., by Efromovich (1999), Härdle (1990), Härdle et al. (2004), Horowitz (1998), Pagan and Ullah (1999), and Ramsay and Silverman (1997).

Local volatility models or their stochastic ramifications are not the only way to price derivatives. Of same significance are approaches relying on stochastic volatility specifications and on Lévy processes. Indeed, the current literature on derivatives pricing may be divided into two main camps: the partisans of local volatility models who prefer them, because local volatility models produce an almost excellent fit to the observed option data; and those

who criticize local volatility models principally for predicting the wrong smile dynamics. It is this second camp that favors stochastic volatility specifications and Lévy models. In this book, we enter the particulars of this debate, but topics like stochastic volatility and Lévy models are only briefly touched. In doing so, we do not intend to argue that these competing modeling approaches are not justified: they certainly are, and there are very good arguments in favor of them. Rather it is our intention to bring together this important strand of literature and to discuss advantages and potential drawbacks. The pricing of derivatives in stochastic volatility models can be found in the excellent textbooks by Fouque et al. (2000) and Lewis (2000), and an outstanding treatment of jump diffusions is provided in Cont and Tankov (2004), or in Schoutens (2003).

Many computations for this book were done in XploRe. XploRe is a software which provides a combination of classical and modern statistical procedures together with sophisticated, interactive graphics. XploRe also allows for web-based computing services. Therefore this text is offered as an e-book, i.e. it is designed as an interactive document with links to other features. The e-book may be downloaded from www.xplore-stat.de using the license key given on the last page of this book. The e-book design offers a PDF and HTML file with links to MD*Tech computing servers.

Organization of the Book

In Chap. 2, we give an introduction into the classical BS model. The basic option valuation techniques are presented to derive the celebrated BS pricing formula. Next, the concepts of IV and the IVS are introduced. Given the model's inconsistency with the empirical evidence, potential directions of relaxing the rigid assumptions are discussed. This will lead to new interpretations of IV as *averages of volatility*. We proceed in discussing the consequences that arise for pricing and hedging in the presence of the smile. A short summary of the literature that investigates IV as a predictor of realized volatility follows. The chapter concludes by giving an account of the potential reasons for the existence of a non-constant smile function.

Chapter 3 is devoted to local volatility. Up to now, the theoretical relationship between implied and local volatility – and finally instantaneous volatility as the measure of the contemporaneous asset price variability – is not as clear-cut as one might wish. In certain boundary situations or asymptotic regimes only has it been possible to make the relation more precise. Figure 1.2 gives an overview of the current state of research. All relations are developed in the course of the next two chapters. The relationship, possibly most important from a practical point of view, is presented by the dotted line, linking implied and local volatility. It represents the so called *IV counterpart of the Dupire formula*, which enables the pricing of exotic options directly from an estimate of the IVS and its derivatives. The chapter discusses several methods to extract local volatility, especially implied tree techniques. Implied trees can be

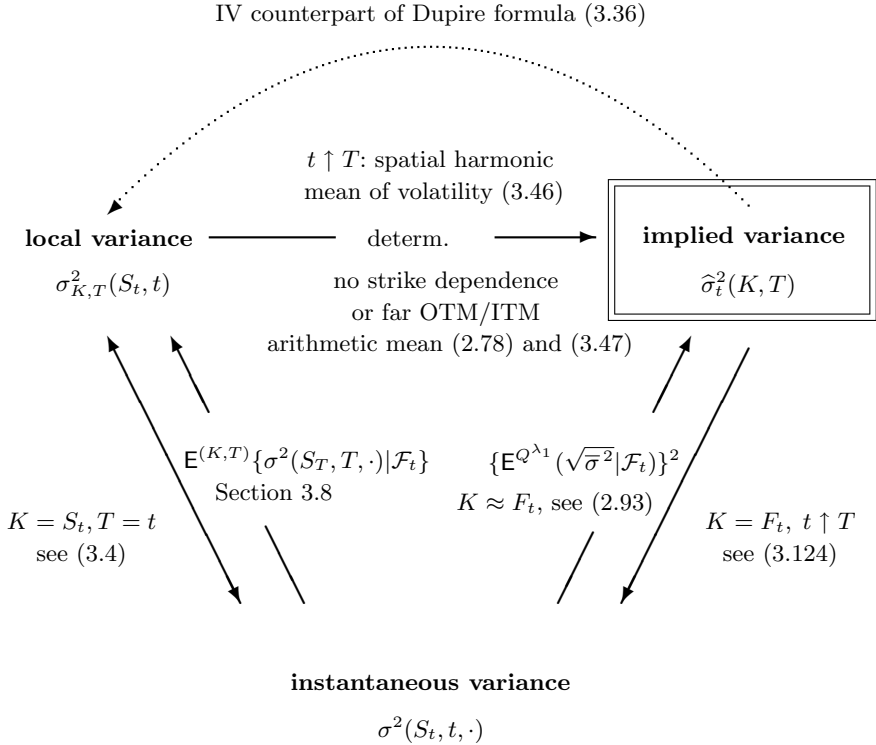


Fig. 1.2. Overview on the volatility concepts important to this work. *Solid lines* denote exact concepts about how the different types of volatility are linked. The *dotted line* represents an ad-hoc relationship. The *arrows* denote the direction of the relation. The term volatility is reserved for objects of the kind σ and $\hat{\sigma}$, while their squared counterparts σ^2 and $\hat{\sigma}^2$ are called variance

considered as nonparametric approximations to the local volatility function. The so called delta debate of local volatility models is covered. The chapter concludes by presenting the class of stochastic IV models.

In Chap. 4, we move to smoothing techniques of the IVS. We introduce the Nadaraya-Watson estimator as the simplest nonparametric estimator for the IVS. This is followed by local polynomial estimation, which is decisive when it comes to the estimation of derivatives. Finally, we introduce a least squares kernel estimator of the IVS. The least squares kernel estimator smoothes the IVS in the space of option prices and avoids the potentially undesirable two-step procedure of previous estimators: traditionally, in the first step, implied volatilities are derived. In the second step the actual fitting algorithm is applied. A two-step estimator may be less biased, when option prices or other input parameters can be observed with errors, only.

The probably biggest challenge in IVS modeling is dimension reduction. This is the topic of Chap. 5, which is divided into two major parts. The first

part, focusses on linear transformations of the IVS. A standard approach in statistics is to apply principal component analysis. In principal component analysis the high-dimensional variables are projected into a lower dimensional space such that as little information as possible is lost. However, this approach is not directly applicable to the IVS due to the surface structure. Hence, we use the *common principal component models* that we find to allow for a parsimonious, yet flexible model choice. A concern of applying the principal component transformation is stability across time. We derive and apply stability tests across different annual samples. The first part concludes by modeling the resulting factors via standard GARCH time series techniques.

The second part of Chap. 5 is devoted to nonlinear transformations via functional principal component techniques. We first outline the functional principal component framework. Then we propose a semiparametric factor model for the IVS. The semiparametric factor model provides a number of advantages compared with other methods: first, surface estimation *and* dimension reduction can be achieved in one single step. Second, it estimates in the local neighborhood of the design points of the surface, only. With regard to Fig. 1.1 this means that we estimate only in the local vicinity of the black dots. This will avoid model biases. Third, the technique delivers a small set of functions and factor loadings that span the propagation of the IVS through space and time. We provide another time series analysis of these factors based on vector autoregressive models and perform a horse race which compares the model against a simpler practitioners' model.

Chapter 6 concludes and gives directions to future research.

The Implied Volatility Surface

A smiley implied volatility is the wrong number to put in the wrong formula to obtain the right price.

Rebonato (1999)

2.1 The Black-Scholes Model

The option pricing model developed by Black and Scholes (1973) and further extended by Merton (1973) is a landmark in financial theory. It laid the foundations of preference-free valuation of contingent claims. Despite its rather restrictive assumptions and the large number of refinements to the model available today, it remains an important benchmark and cornerstone of financial model building. Here, we give a short review of the BS model and present the fundamental results necessary for the further development of this work. For a more detailed account, we refer to textbooks in Finance, such as Musiela and Rutkowski (1997) or Karatzas (1997).

We consider a continuous-time economy with a trading interval $[0, T^*]$, where $T^* > 0$. It is assumed that trading can take place continuously, that there are no differences between lending and borrowing rates, no taxes and short-sale constraints.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and $(W_t)_{0 \leq t \leq T^*}$ a Brownian motion (see appendix Chap. B for a definition of the Brownian motion) defined on this space. \mathbb{P} is the objective probability measure. Information in the economy is revealed by a filtration $(\mathcal{F}_t)_{0 \leq t \leq T^*}$, which is the \mathbb{P} -augmentation of the natural filtration

$$\mathcal{F}_t^W = \sigma(W_s, 0 \leq s \leq t), \quad 0 \leq t \leq T^*. \quad (2.1)$$

The filtration is assumed to satisfy the ‘usual’ conditions, namely that it is right-continuous, and that \mathcal{F}_0 contains all null sets.

The asset price $(S_t)_{0 \leq t \leq T^*}$, which pays a constant dividend yield δ , is modelled by a geometric Brownian motion adapted to $(\mathcal{F}_t)_{0 \leq t \leq T^*}$. The evolution of the asset is given by the stochastic differential equation (SDE):

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t, \quad (2.2)$$

where μ denotes the (constant) instantaneous drift and σ the (constant) instantaneous (or spot) volatility function. The quantity σ^2 measures the instantaneous variance of the return process of $\ln S_t$. Thus, instantaneous volatility