
Option Valuation under Stochastic Volatility

With Mathematica Code

Copyright © 2000 by Alan L. Lewis

All rights reserved. Except for the quotation of short passages for the purposes of criticism and review, no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher.

Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or the consequences of their use. All information, including formulas, documentation, computer algorithms, and computer code are provided with no warranty of any kind, express or implied. Neither the author nor the publisher accept any responsibility or liability for the consequences of using them, and do not claim that they serve any particular purpose or are free from errors.

Published by: Finance Press, Newport Beach, California, USA
Cover design by: Brian Burton Design

For the latest available updates and corrections to this book and publisher contact information, visit the Internet sites:
<http://www.financepress.com> or <http://members.home.net/financepress>

Trademark notice: *Mathematica* is a trademark of Wolfram Research, Inc., which is not associated with the author or publisher of this book.

International Standard Book Number 0-9676372-0-1
Library of Congress Card Number: 99-91935

Printed in the United States of America
Printed on acid-free paper

1 Introduction and Summary of Results

Suppose we use the standard deviation ... of possible future returns on a stock ... as a measure of its volatility. Is it reasonable to take that volatility as constant over time? I think not.

Fischer Black

The problem of how to value an option is a fascinating one, with a relatively long history for a financial topic. By 1900, Bachelier suggested a fair game approach using a normal distribution for the underlying security price—almost the modern model. But Bachelier’s approach became suspect once the ideas of utility theory and risk-aversion became central in economic theory. If most investors were averse to risk-taking, shouldn’t that influence an option’s price? Shouldn’t an option be worth less than its fair game value, just like a stock?

This puzzle was resolved by Black and Scholes (B-S 1973) for a world where security prices followed a geometric Brownian motion process with constant volatility. Their elegant solution for option pricing is, surprisingly, completely independent of investors’ risk attitudes and expectations.

However, the assumption of constant volatility was suspect from the beginning, as the 1976 quotation above suggests. The facing figure shows the monthly volatility realized by the S&P 500 Index returns from 1928-1999. Statistical tests strongly reject the idea that a constant volatility process could have generated these data. It also became clear, although this was less immediate, that the B-S model was in conflict with evolving patterns in observed option pricing data. In particular, after the 1987 market crash, a persistent pattern emerged, called the “smile” that shouldn’t exist under the B-S theory.

In this book we study option valuation when security prices evolve with stochastic (random) volatility. Stochastic volatility models lead to generalizations of the B-S option pricing formula. The generalized models are both mathematically interesting and useful because they can explain the real-world patterns that are missing from the B-S theory. Not only do stochastic volatility models explain the basic shapes of smile patterns, but they also allow for more realistic theories of the “term structure” of implied volatility. Consequently, the material should be of interest to both financial academics and traders.

Option valuation is still an extremely active area in finance, with many references cited in the notes. We don’t create a comprehensive survey—instead, our goal is to create a uniform and fairly comprehensive theoretical treatment of the case where the stock price and volatility are described by a two-dimensional diffusion process. We work exclusively in continuous-time with continuous sample paths.

Once you allow random volatility into the theory, utility functions and risk attitudes come back into the theory also. In a B-S world, where volatility is constant, option prices are determined purely by arbitrage arguments and risk attitudes play only an indirect role. (That is, to set the price of the primary securities). But, with stochastic volatility, the absence of arbitrage alone does not fix option prices. We discuss this in detail in Chapter 7 and summarize our results below.

We consider a number of particular stochastic processes for volatility. However, we have a particular interest in the case where volatility is described as the diffusion limit of a GARCH-type process—a GARCH diffusion, for short. Since the seminal work of Engle (1982), discrete-time ARCH models have become a proven approach to modeling security price volatility. For a review of the substantial literature, see Bollerslev, Chou, and Kroner (1992). Since Nelson (1990), it has been understood that GARCH-type models have well-defined continuous-time limits.

From a practical point of view, the advantage of the GARCH diffusion model is that you can estimate its parameters by using off-the-shelf software for GARCH processes. (See Appendix 1 to this chapter). Although no closed-form solution is

available for options priced under the particular process we call a GARCH diffusion, we are able, nevertheless, to develop a fairly complete picture.

1 Summary of Results

Our security model for most of this book is an (equity) price process P of the general form

$$(1.1) \quad P : \begin{cases} dS_t = (\alpha_t S_t - D_t)dt + \sigma_t S_t dB_t \\ dV_t = b(V_t)dt + \xi \eta(V_t) dZ_t \end{cases},$$

where dB_t and dZ_t are Brownian motion processes with correlation $\rho(V_t)$. In (1.1) α_t is the instantaneous expected total return of the stock, which pays the owner dividends at the dollar rate D_t . The volatility $V_t = \sigma_t^2 \geq 0$, where σ_t is the instantaneous standard deviation of the price returns. Of course, we usually call σ_t the volatility also. In addition to using P as a simple label, as in (1.1), we will sometimes use P to mean a probability (measure): namely, the one under which dB_t and dZ_t are Brownian motions. That alternative interpretation should cause no confusion.

The “volatility of volatility” parameter ξ sets the scale for the random nature of volatility and plays a very important role in the theory. The functions $b(V)$ and $a(V) = \xi \eta(V)$ are the drift and diffusion coefficients of the (actual) volatility process. We consider only actual volatility processes that are time-homogeneous. Sometimes the risk-adjusted volatility process will pick up a time dependence from the risk-adjustment.

Examples of stochastic volatility models

I. The GARCH Diffusion Model. The GARCH diffusion model is important because it’s the continuous-time limit of many GARCH-type processes.¹ In this book, we use ‘GARCH diffusion’ to mean an actual volatility process of the form:

¹ In fact, the term GARCH is a loose term that accommodates many types of discrete-time financial models and various continuous-time limits are possible. For example, Heston and Nandi (1997) have recently shown that a degenerate case of Heston’s (1993) square root model (see below) can be obtained as a limit of a particular GARCH-type process.

$$(1.2) \quad dV_t = (\omega - \theta V_t) dt + \xi V_t dZ_t .$$

The volatility drift parameters, ω and θ , are assumed to be constants and capture the mean-reverting nature of the volatility process. Since θ has the dimensions of inverse time, $1/\theta$ represents a “half-life” for volatility shocks.

Empirically, estimates for $1/\theta$ are quite variable against stock indices and can range from a few weeks to more than a year. For U.S. broad-based stock indices, ξ is typically in the range of one to two on an annualized basis, which represent volatility uncertainty of 100 to 200% over a year. The correlation, ρ , captures the association between security price changes and volatility changes. Typically, negative price shocks are associated with higher volatility than positive shocks of the same magnitude. For the same indices, we find $\rho \cong -0.5$ to -0.8 .

See Appendix 1 of this chapter for a simple discrete-time GARCH model that admits (1.2) as a continuous-time limit. In that Appendix, we provide more examples of parameter estimates and formulas for estimators.

II. The Square Root Model. The square root model is described by the volatility process

$$(1.3) \quad dV_t = (\omega - \theta V_t) dt + \xi \sqrt{V_t} dZ_t .$$

This model is very important because of two facts. First, as shown by Heston (1993), it has essentially a closed-form solution for option prices which is easy to implement. Second, that solution is *typical*: it displays the same qualitative properties that we expect in general time-homogenous cases. Consequently, it tells us, for example, how the (unknown) solution to the GARCH diffusion behaves in many respects.

For the square root model, clearly ω and θ play the same role as in the GARCH diffusion. Also ξ again sets the scale for the volatility of volatility, but its numerical magnitude will be quite different when estimated against the same pricing data because it now multiplies a factor of $\sqrt{V_t}$ instead of the factor of V_t in the GARCH case. From dimensional considerations, $\xi_{1/2} \approx \xi_1 \bar{\sigma}_t$, where $\xi_{1/2}$ is the square root model parameter, ξ_1 is the GARCH diffusion parameter, and $\bar{\sigma}_t$ is the average volatility of the data series being considered. For example, for the S&P 500 Index, if ξ_1 is in the range (1,2) annualized, and $\bar{\sigma}_t \approx 10\% = 0.1$ annualized, then $\xi_{1/2}$ is in the approximate range (0.1, 0.2) annualized.

III. The 3/2 model. The 3/2 model is described by the volatility process

$$(1.4) \quad dV_t = (\omega V_t - \theta V_t^2) dt + \xi V_t^{3/2} dZ_t.$$

This model is important because, not only does it have a closed-form solution almost as simple as the square root model, but it displays a feature of many stochastic volatility models that you don't see in the square root model. That is, even after a change of measure to the risk-adjusted process, option prices (relative to the bond price) under the 3/2 model are sometimes not martingales, but merely local martingales. When option prices are not martingales, this means that they are not given by the standard expected value formula—for example, $e^{-rT} \mathbb{E}_t[(S_T - K)^+]$ for a call option. Here \mathbb{E}_t denotes an expectation under the pricing process (see the subsection below titled “the risk-adjusted process”).

This failure of the usual martingale pricing relation is actually quite common in (unbounded) stochastic volatility models under very typical risk-adjustments, such as log-utility. The failure can also occur in the GARCH diffusion, for example, but only a specialized (and very complicated) closed-form solution is currently available for that one. So the 3/2 model is one of the simplest illustrations of this important phenomenon for financial theory, and our results for that model illustrating this behavior, mostly discussed in Chapter 9, are new.²

Again, from dimensionality, one expects $\xi_{3/2} \cong \xi_1 (\bar{\sigma}_t)^{-1}$. In this case, for the S&P 500 Index, if ξ_1 is in the range (1,2), and $\bar{\sigma}_t \cong 0.1$, then $\xi_{3/2}$ is in the approximate range (10, 20), again with all values annualized.

IV. The Ornstein-Uhlenbeck process. With $y_t = \ln V_t$, this model is described by the process

² Other aspects of the 3/2 model have been independently developed by Heston (1997). The failure of the martingale formula for the stock price in the GARCH diffusion model was first shown by Sin (1998). These failures are specific examples of the notion that the absence of arbitrage implies that financial claim prices are, *in general*, only strictly local martingales—not martingales. This refinement of earlier notions about financial claim prices as martingales is due to Delbaen and Schachermayer (1994).

$$dy_t = (\tilde{\omega} - \theta y_t) + \xi dZ_t, \quad (-\infty < y_t < \infty)$$

By using the Ito formula (explained in Sec. 2), and letting $\omega = \tilde{\omega} + \frac{1}{2}\xi^2$, the process for V_t is

$$dV_t = (\omega V_t - \theta V_t \ln V_t) dt + \xi V_t dZ_t.$$

We don't treat this model specifically, but it's a very popular one among researchers. Many of our general results apply to it.

The risk-adjusted process. To value an option, you don't use (1.1), but a closely related process \tilde{P} which is often called the risk-adjusted process. Throughout this book, it has a form similar to (1.1):

$$(1.5) \quad \tilde{P}: \begin{cases} dS_t = (rS_t - D_t)dt + \sigma_t S_t d\tilde{B}_t \\ dV_t = \tilde{b}(V_t, t)dt + \xi \eta(V_t) d\tilde{Z}_t \end{cases}$$

To get from (1.1) to (1.5) we have made two changes: (i) we replaced the equity expected return by an interest rate r and (ii) we replaced the volatility drift by another function $\tilde{b}(V_t, t)$, the risk-adjusted volatility drift. This procedure is carried out explicitly for a class of equilibrium models in Chapter 7: representative agent models with power utility functions. In addition, we will often assume that the representative agents have very distant planning horizons.³ This allows one to further simplify (1.5) to a volatility process with a time-homogenous drift $\tilde{b}(V_t)$.

Equation (1.5) has an associated partial differential equation (PDE) that determines option prices (see below). In the special case where the dividend yield is constant, then (1.5) is proportional in the stock price and the PDE may be solved with a transform technique.

The transform idea is explained in Chapter 2, where we introduce the *fundamental transform* $\hat{H}(k, V, \tau)$. Besides the volatility V , the transform depends upon k , the (generalized) Fourier transform variable and τ , the time to the option's expiration. The fundamental transform is determined by the volatility process and not by the particulars of any option contract. Once you have it, the *same* function $\hat{H}(k, V, \tau)$ is used to determine the value of *every*

³ The planning horizon is the "date of death", not to be confused with the option expiration date, which may be very near.

European-style financial claim. This last step only requires an integration in the complex k -plane. So you can see why this function deserves its title.

While the idea of a transform-based approach is not new, previous applications have tended to be model-specific. Not only are our results more general, but they encompass the situation when option prices, relative to a numeraire, are not martingales, but only strictly local martingales. Specifically, we develop two formulas for option prices, which we call Solution I and Solution II. Solution I is the usual martingale-style or expected value formula. Solution II is more general and sometimes includes an additional term that relates to volatility “explosions”. When a call option price is not a martingale, the Solution II formula yields the desired arbitrage-free fair value.

The fundamental transform has a lot of special properties, many of which follow from the fact that not only is it a characteristic function, but it’s an *analytic characteristic function*. In particular, the function is regular in strips in the complex k -plane and also has the *ridge property*. The ridge property, which is more or less what it sounds like, plays a very important role in the theory of the term structure of implied volatility. This property is explained in Chapter 2 and applied in Chapter 6.

In Chapter 3, we show to solve the valuation PDE associated with (1.5) by a power series in ξ , the volatility of volatility. The first two terms of that series, which can be rapidly evaluated on any desk-top computer, may suffice for most trading applications.

In Chapter 4, we discuss “mixing theorems”, which show that solutions to the valuation PDE are weighted sums of the B-S solutions. This important concept has a number of applications, including a new Monte Carlo method.

In Chapter 5, we take up the theory of the smile. The volatility of volatility series expansion generates an explicit formula for the smile through order ξ^2 that explains the qualitative patterns seen in the marketplace. Also we explore the accuracy of “quadratic approximations” to the smile by developing correction terms.

The term structure of implied volatility is the subject of Chapter 6. One new result is that we show that the asymptotic theory is very simple: the asymptotic

implied volatility is given by $V_{\infty}^{imp} = 8\lambda(k_0)$, where λ is the first eigenvalue of a differential operator and k_0 is a complex number. Regardless of the current volatility and regardless of the degree of “moneyness”, the smile flattens as a function of the time to expiration to this same value V_{∞}^{imp} .

In Chapter 7, we present a utility-based equilibrium theory. This theory, with a distant planning horizon, provides several new and explicit examples of transformations of the volatility drift function $b(V)$ into a risk-adjusted drift $\tilde{b}(V)$. The model is the representative agent theory. The representative has a power utility function with a constant proportional risk-aversion (CPRA) parameter γ , where $\gamma = 1$ corresponds to risk-neutrality. Under this CPRA equilibrium, we find a simple pattern for the effect of risk attitudes on option prices: with $\rho \leq 0$, call or put option prices are raised in an interval $\gamma_1(\rho) < \gamma < 1$ and lowered in an interval $\gamma < \gamma_1(\rho)$ relative to their risk-neutral values.

Another key, and novel result from Chapter 7 is the following. When the representative agent is a pure investor (no consumption until a distant planning horizon), then the risk-adjustment problem has a simple solution. The volatility risk premium (defined below in Sec. 5) is given by

$$\lambda^v(V) = (1 - \gamma)\rho(V)\sqrt{V} - a(V)\psi(V), \quad \text{where} \quad \psi(V) = u'(V)/u(V).$$

And $u(V)$ is the first eigenfunction of the *same* differential operator that determines λ above. The eigenfunction depends upon γ . For a pure investor with any planning horizon, and $\gamma < 1$, we prove that when $\psi(V, t)$ exists, it has the same sign as γ . The leading eigenfunction is easily computed in Mathematica in just a few seconds.

In Chapter 8, we discuss the change-of-numeraire transformation or duality for short. This transformation generalizes a known symmetry (put-call symmetry) under constant volatility. Under duality, two representative agents who are risk-neutral and have log-utility, respectively, switch places. The general effect of the change-of-numeraire transformation on stochastic volatility models has been previously noted by Schroder (1999) and is implicit in Sin (1998); what is new here is placing the transformation in the context of our CPRA equilibrium.

In Chapter 9, we give a detailed account of the effect of volatility explosions and the failure of the martingale pricing formula. In economically reasonable

models, the actual volatility is recurrent and, in many models, stationary. But, after risk-adjustment either the risk-adjusted volatility, or a closely related process that we call the auxiliary volatility process, can explode. If the risk-adjusted volatility process is written $dV = \tilde{b}dt + a d\tilde{Z}$, then the auxiliary volatility process is $dV = (\tilde{b} + \rho a \sigma)dt + a d\tilde{Z}$. An explosion means that the volatility can reach plus infinity in finite expected time. Our main result is that, when the auxiliary volatility process can explode prior to τ , with probability $\hat{P}_e(V, \tau)$, then the call option price is not a martingale. Instead, the price is given by

$$C(S, V, \tau) = e^{-r\tau} \mathbb{E}_t[(S_T - K)^+] + Se^{-\delta\tau} \hat{P}_e(V, \tau).$$

We also show that this formula is identical to the Solution II formula of Chapter 2. These results, which are new, are also extended to any “call-like” claim.

In Chapter 10, we study both the fundamental transform and option prices at large volatility. This behavior is important in a number of contexts in previous chapters. We find that, in contrast to the B-S model, it is not always true that the call option price tends to the stock price as the volatility increases to infinity. A specific counter-example is found.

Finally, in Chapter 11, we develop the closed-form solutions for the fundamental transform in our three running examples: the square root model, 3/2 model, and a special case of the GARCH diffusion (geometric Brownian motion).

Mathematica code. For readers who are Mathematica users, many code fragments are included, usually in Appendices. All of the Mathematica routines can be evaluated in short periods, ranging from a few seconds to a few minutes on a desk-top machine. But the book can also be read without these sections. No attempt is made to explain the Mathematica system or its built-in objects, but I do attempt to explain how the code fragments work in a general sense.

Notations. Frequently used notations are given on a page at the end of the book. I refer to equation numbers within a chapter by their label within that chapter: say (2.3) for the third equation in Sec. 2 of Chapter 7. But, in another chapter, I refer to that same equation as (7.2.3). The more important equations are placed in boxes for emphasis.

Subjects not covered

Discrete-time theory. Virtually all of the development in this book can be translated into a discrete-time setting. But space and time limitations, for one thing, have forced me to omit this. Also, because the results do readily translate, the discrete-time setting for the option valuation equations don't really tell you anything new.

Having said that, there is an important caveat. The continuous-time setting used here, with its continuous sample paths may not be "variable enough" for the real world. The discrete-time world, especially GARCH models, accommodates "wider than normal distributions" in a straightforward way. So this may lead one back to a discrete-time theory.

One-dimensional models. There is a large literature that considers one-dimensional diffusions of the form $dS_t = rS_t dt + \kappa(S_t)dB_t$ as an explanation for smile and term structure patterns. These models are difficult to interpret over long time periods in a world with exponentially growing stock prices—so, for the most part, we don't discuss them. The empirical performance of this type of model is reviewed by Dumas, Fleming, and Whaley (1998). One exception is that, in Chapter 9, we study $dS = S^\varphi dB_t$, $\varphi > 1$ as a simple example of the failure of the martingale pricing formula.

Empirical comparisons. This important area deserves a book in itself, and by those more qualified. One very nice study, which is important to the models here, is a careful examination of put-call parity by Kamara and Miller (1995). All of the theoretical models in this book satisfy the put-call parity relation, which relates European-style put and call values by an arbitrage argument. In practice, there are many small violations, but in investigating all violations in intraday transaction data, Kamara and Miller found that almost half of the "arbitrages" result in a loss when execution delays are accounted for. Moreover, the mean ex post profit in trying to exploit the violations was negative.

American-style options. The effect of stochastic volatility on the early exercise premium for American-style options is an important subject. With one exception, we only treat European-style options. The exception is Chapter 8, which presents the duality relation mentioned above. This relation, as it does in the constant volatility world, connects the early exercise boundaries of the call option and its dual put partner.

The remaining sections in this chapter review some hedging arguments under both constant and stochastic volatility and the martingale pricing argument. This material is mostly standard. Advanced readers might want to look at some comments about risk premiums in Sec. 4, comments about local martingales in Sec. 5, and then skip directly to Chapter 2.