Journal of Banking & Finance 34 (2010) 1982–1992

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Estimating financial risk measures for options

Ghulam Sorwar *, Kevin Dowd

University of Nottingham, UK

ARTICLE INFO

Article history:
Received 17 March 2009
Accepted 7 January 2010
Available online 11 January 2010

JEL classification:
G11
G13
C15

Keywords:
Financial risk measures
Path dependent options
Fat tail
Simulation

ABSTRACT

This paper proposes a simulation-lattice procedure to estimate financial risk measures for option positions. The framework proposed can be applied to many different kinds of options, including exotic and vanilla options; it can take account of early exercise features; heavy tails in underlying processes; estimate different risk measures, including VaR, Expected Shortfall and Spectral Risk Measures; and in a limited way it can be generalized to accommodate multiple-factors. It avoids many of the limitations of existing approaches and, in particular, avoids the problems associated approaches based on delta–gamma and similar approximations. It also generates some interesting results about the risk measures of some illustrative options positions.

1. Introduction

Options form a major category of modern financial assets, and are important risk management tools. However, measuring the risks of option positions is notoriously tricky. This is due to a combination of difficult factors: non-linearities, time-dependence (leading to changes in option Greeks as expiry approaches), underlying instability (leading to unstable vegas, etc.), the need to take account of early exercise features, the problems associated with heavy tails in financial returns, option payoffs do not have closed form distribution at expiry and the complications associated with multiple risk factors. It is therefore, not surprising that attempts to obtain financial risk measures for options positions have met with only limited success.[1] For example, closed-form[2] and algorithmic[3] solutions for options risk measures are few and far between. The most common approaches are those based on delta–gamma and similar approximations, but these can be inaccurate, do not take account of early exercise, and become unwieldy with multiple risk factors.[4] Simulation approaches that do not rely on any approximations have also been suggested, but these have tended to be computationally expensive and also have difficulty taking account of early exercise.[5] Furthermore, to our knowledge none of these approaches takes account of possible heavy tails in the returns of underlying financial assets: heavy tails are important not only because they can have a major impact on option values, but also because it is well-known that ignoring heavy tails in other contexts can often lead to large under-estimates of the risks (and, hence, of the financial risk measures) involved. Consequently, the measurement of options risks is one of and possibly the most difficult area in the whole of modern financial risk measurement.

This paper proposes an approach to options risk measurement that can deal with many of these problems. The basic idea is a combination of simulation and lattice: we use simulation to generate the underlying stock price and a lattice to obtain the values of any unexercised options at the end of the holding period. We take account of early exercise by breaking the stock price path into

* Corresponding author. Tel.: +44 (0) 115 951 5487; fax: +44 (0) 115 846 6667.
E-mail addresses: ghulam.sorwar@nottingham.ac.uk (G. Sorwar), kevin.dowd@hotmail.co.uk (K. Dowd).

[1] We use the term 'financial risk measure' in this paper to refer to risk measures based on the quantiles of the prospective loss distribution for the options considered. Financial risk measures so defined might refer to VaRs, Expected Shortfalls and comparable (e.g. coherent or spectral) risk measures. The Greeks might be indicative of options' risks, but are not financial risk measures by this definition.

[2] There are very few closed-form solutions for options risk measures, and even these are limited to only to the VaR. Several examples are given in Dowd (2005, pp. 249–253). See also Rouvinez (1997).

[3] For example, Fong and Lin (1999) propose an algorithmic approach to obtain the VaRs of various options. However, this approach is a little limited because it depends on local approximations around identifiable minima.

[4] Many different types of delta–gamma approach have been used. For more on their drawbacks, see Britten-Jones and Schaefer (1999), Wiener (1999) or Dowd (2005, pp. 256–264).

[5] For more on these issues, see, e.g., Broadie and Glasserman (1998, pp. 202–205) or Dowd (2005, Chapter 10).

0378-4266/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2010.01.005

discrete intervals and periodically check for early exercise. Where it is optimal to exercise early, we assume that the option is exercised and the proceeds held in cash till the end of the holding period. Thus, our approach is applicable to options with early exercise features (such as American and Bermudan options) as well as European options. Options are valued using a trinomial lattice, and the flexibility of the lattice allows us to handle a great many different types of option such as barrier options, 'best of' options, Asian options, etc., as well as vanilla ones. Heavy-tails are allowed for using a lattice suggested a few years ago by Boyle and Tian (1999): this lattice accommodates tail heaviness by working with a constant elasticity of variance diffusion process rather than the more traditional lognormal diffusion process.[6] This lattice is also easily adapted to handle early exercise features of American options, in contrast to other lattice approaches such as the one proposed by Topaloglou et al. (2008) which is restricted to European options. And unlike some existing approaches to options risk measurement, our approach allows us to generate a range of different financial risk measures: besides enabling us to obtain the VaR, it also enables us to obtain more recently developed (and superior) financial risk measures such as the Expected Shortfall (ES) and Spectral Risk Measures (SRMs) as well.

The paper is organized as follows. Section 2 sets out the option valuation (i.e. lattice) approach to be used for underlying asset following a CEV process. This approach can be altered to another stochastic process with appropriate adjustments. Section 3 explains the different risk measures to be obtained: the VaR, the ES and the SRM. Section 4 outlines the simulation framework used to obtain our financial risk measures and making use of the lattice for valuation purposes. Section 5 presents some results for illustrative options. Section 6 concludes.

2. Option valuation

Our first requirement is a model to value options. A natural starting point is to assume that the underlying stock price follows a geometric Brownian motion:

$$dS = \mu S\,dt + \sigma S\,dz, \tag{1}$$

where $z(t)$ is standard Brownian motion. The percentage change in the stock price over a small interval dt is then normally distributed with mean $\mu\,dt$ and variance $\sigma^2\,dt$. However, it is well-known that this process provides an empirically poor fit to stock prices.

We can get a better fit if we follow Cox (1996) and Cox and Ross (1976) and work with the following constant elasticity of variance (CEV) diffusion process:

$$dS = \mu S\,dt + \sigma S^{\frac{\alpha}{2}}\,dz, \tag{2}$$

where $0 \leqslant \alpha \leqslant 2$. If $\alpha < 2$, the process exhibits heavier tails than (1). This means that tails become heavier as α falls, but (2) reduces to the lognormal case (1) in the limiting case when $\alpha = 2$. Thus, (2) is a more general process than (1) and encompasses the latter as a special case. From an empirical perspective, (2) is superior to (1) because it can accommodate the heavier-than-normal tails commonly observed for stock returns.[7]

Eq. (2) gives us the real or actual price process (or P-measure). To value derivatives, we need the risk-neutral (or Q-measure), and we obtain this in the usual way by replacing μ with $r - q$, where we assume that our stock yields a continuous dividend yield q, viz.:

$$dS = (r-q)S\,dt + \sigma S^{\frac{\alpha}{2}}\,dz. \tag{3}$$

Derivatives can now be priced by taking the expectation of the payoff at maturity under the Q-measure.

We would now usually price the options using a trinomial lattice.[8] However, applying a lattice directly to (3) would lead a major problem: the presence of the $S^{\frac{\alpha}{2}}$ term creates a heteroskedasticity that means that our lattice does not recombine, and a non-recombining lattice would be computationally extremely expensive. To get around this problem, we use the trinomial lattice proposed by Boyle and Tian (1999). This is based on the idea that we first transform S so that the heteroskedasticity disappears; we then construct a recombining lattice on the transformed variable, and use this to recover the option value. To implement this lattice, we first transform the stock price S to a new process $y(t,S)$ that has constant volatility. Using Ito's lemma

$$dy = q(t,S)\,dt + \frac{\partial y}{\partial S}\sigma S^{\frac{\alpha}{2}}\,dz, \tag{4}$$

where

$$q(t,S) = \frac{\partial y}{\partial t} + (r-q)S\frac{\partial y}{\partial S} + \frac{1}{2}\sigma^2 S^{\alpha}\frac{\partial y^2}{\partial S^2}.$$

To ensure constant volatility, we find a transformation such that

$$\frac{\partial y}{\partial S}\sigma S^{\frac{\alpha}{2}} = v \tag{5}$$

for some constant v. Eq. (5) gives us

$$\frac{\partial y}{\partial S} = \frac{v}{\sigma}S^{\frac{-\alpha}{2}} \tag{6}$$

and we integrate this to obtain

$$y = \begin{cases} \frac{v}{\sigma(1-\frac{\alpha}{2})}S^{1-\frac{\alpha}{2}} \\ \frac{v}{\sigma}\ln S \end{cases} \quad \text{for} \quad \begin{cases} \text{for } \alpha \neq 2 \\ \text{for } \alpha = 2 \end{cases}. \tag{7}$$

For the general case $\alpha \neq 2$, the transformed process is

$$dy = \left[r\Big(1-\frac{\alpha}{2}\Big)y - \frac{\alpha v^2}{4(1-\frac{\alpha}{2})y}\right]dt + v\,dz \tag{8}$$

(8) has the desired constant volatility. However, as Boyle and Tian explain, this process also has a drift term which generally explodes when $y \to 0$, and this exploding drift causes problems for standard trinomial lattices, because the jumps and probabilities must be chosen to match the drift as well as the volatility.

We can get around this problem using a suggestion first made by Tian (1994). He suggests that the transformed process (8) be used to define an evenly spaced two-dimensional grid in the (t, y) space represented by the partitions

$$t_0, t_0+\Delta t, t_0+2\Delta t, \ldots, t_0+N\Delta t$$
$$y_0 - M_1\Delta y, \ldots, y_0-\Delta y, y_0, y_0+\Delta y, y_0+M_2\Delta y$$

where the trinomial process starts at (t_0, y_0) and Δt and Δy are the increments in time and y. To ensure stability, the y step is chosen as

[6] Our approach is also capable of some limited extension to multi-factor cases. However, as things currently stand, it is difficult to allow for both multi-factors as well as tail heaviness.

[7] In addition, the CEV process also has the advantage that it links the volatility to the level of the stock price, and can so accommodate the empirical stylized fact that volatility tends to rise as the stock price falls. As Cox (1996) suggested, it is possible that this negative correlation may be the source of the volatility smile usually encountered using Black-Scholes.

[8] We would typically use a lattice procedure for non-vanilla options, but an additional reason for using a lattice is to accommodate early exercise features.