International Review of Financial Analysis xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

International Review of Financial Analysis

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

Test of recent advances in extracting information from option prices

J.V. Healy[a], A. Gregoriou[a], R. Hudson[b,*]

[a] Brighton Business School, University of Brighton, United Kingdom
[b] Hull University Business School, United Kingdom

ARTICLE INFO

Keywords:
Option pricing
Risk neutral moments
Risk neutral density
Analytical solutions

JEL classifications:
C14
G11
G12

ABSTRACT

A large literature exists on techniques for extracting probability distributions for future asset prices from option prices. No definitive method has been developed however. The parametric 'mixture of normals', and non-parametric 'smoothed implied volatility' methods remain the most widespread approaches. These though are subject to estimation errors due to discretization, truncation, and noise. Recently, several authors have derived 'model free' formulae for computing the moments of the risk neutral density (RND) directly from option prices, without first estimating the full density. The accuracy of these formulae is studied here for the first time. The Black-Scholes formula is used to generate option prices, and error curves for the first 4 moments of the RND are computed using the 'model-free' formulae. It is found that, in practice, the formulae are prone to large and economically significant errors, because they contain definite integrals that can only be solved numerically. We show that without mathematically equivalent expressions with analytical solutions the formulae are difficult to deploy effectively in practice.

1. Introduction

Methods for extracting implied probability distributions for the prices or returns of an asset at a future time, from series of synchronously observed market prices of options on the asset, have been extensively studied since the mid 1990s [see e.g. Bliss and Panigirtzoglou (2002), Jackwerth (2004), and more recently Figlewski (2008) for reviews]. Breeden and Litzenberger (1978) made explicit the exact relationship between option prices and the risk neutral density (RND) [see Appendix A for details and proof]. In the risk-neutral pricing framework the price of an option is equal to its discounted expected payoff under the risk neutral measure. Evaluating the integral of the payoff function over the risk neutral measure and discounting at the risk free rate can thus price an option. Given a continuum of observed option prices, this pricing calculation can be inverted for European exercise options, and the full RND for the price (return) of the underlying asset at maturity extracted. Useful information contained in the shape of the distribution can thus be recovered. RNDs have numerous important applications in finance. These include: Pricing securities [Cox & Ross, 1976]; Estimating value-at-risk (VaR) for risk management purposes [Ait-Sahalia & Lo, 2000]; Studying risk aversion and risk preferences [Bliss & Panigirtzoglou, 2004]; Assessing financial market expectations regarding future asset prices, interest rates, and exchange rates, in connection with setting monetary policy [Lynch & Panigirtzoglou, 2008]. However, existing methods for extracting RNDs are variously, computationally cumbersome, data intensive, and or subject to estimation errors due to discreteization, truncation, and noise issues in the raw options data. No definitive method has been developed, but two approaches are popular with practitioners, namely, the mixture of normals [Ritchey, 1990], and the smoothed implied volatility method [Shimko, 1993]. Tests suggest the latter method produces better results [see e.g. Bliss & Panigirtzoglou, 2002 and Andersson & Lomakka, 2003].

In many applications it is enough to know the first four moments of the RND. Hence a more parsimonious representation will suffice. Based on recent theoretical developments, several authors have derived exact formulae for computing the moments of the RND directly from option prices without first estimating the full density distribution. These formulae have the advantage of being 'model free', in the sense of not being subject to the assumptions of any option pricing model. New approaches for extracting the RND, by using these formulae to compute its moments in a first step, have also been developed. Of course, when these formulae are applied to observational data, they also are subject to estimation errors due to discretization, truncation, and noise issues in the data. Jiang and Tian (2007), and Dennis and Mayhew (2002, 2009), have studied the errors arising from discrete implementation of the 'model free' implied variance, and the implied skewness and kurtosis respectively, and show that they are economically significant. What is perhaps less well appreciated is that the solutions to the formulae themselves exhibit sensitivity to their inputs, even for realistic ranges of values, and are thus capable of being biased estimates.

* Corresponding author at: Hull University Business School, University of Hull, Hull HU6 7RX, United Kingdom.
E-mail address: robert.hudson@hull.ac.uk (R. Hudson).

http://dx.doi.org/10.1016/j.irfa.2017.09.011
Received 30 January 2017; Received in revised form 22 September 2017; Accepted 24 September 2017

The contribution of this paper is to demonstrate that applying the 'model free' formulae for the first four moments of the RND produce large and economically significant errors independently of those resulting from the observation issues discussed above. This is shown by solving the formulae as exactly as possible in a continuous strike price framework, for realistic ranges of inputs and constructing error curves. This is important because it shows that the formulae are of limited applicability in their current forms.

The remainder of the paper is organised as follows. Section 2 reviews the literature on 'model free' implied moments of the RND. Section 3 outlines the methodology used. Section 4 presents the findings. Section 5 contains a summary and conclusions.

2. 'Model free' implied moments: literature review

The development of "model-free" methods of directly extracting the moments of the RND has emerged from three separate strands of research. First, work on the log contract, volatility, and variance swaps [Neuberger (1994), Carr and Madan (1998), Demeterfi, Derman, Kamal, and Zou (1999)]. Second, extraction of information on the underlying price processes from option prices [Derman & Kani, 1994; Britten-Jones & Neuberger, 2000]. Third, studies of the characteristic function of the state price density (discounted RND), as an alternative spanning entity to options for pricing other securities [Bakshi & Madan, 2000; Bakshi, Kapadia, & Madan, 2003].

Demeterfi et al. (1999) show how hedging an option on the logarithm of the price of an underlying asset (the log contract), provides a payoff equal to the variance of the asset's returns. No such contract is traded in practice; however the log contract can be replicated by a portfolio of European exercise options with a continuous range of strikes and maturities. This portfolio has a value equal to the payoff of the log contract. Dynamically hedging a log contract therefore captures realized variance (volatility). The value of a variance swap, a forward contract F on future realized variance with strike K, depends on the future payoff $({\sigma_R}^2 - K_{VAR}) \times N$ discounted to its present value under the risk neutral measure, where ${\sigma_R}^2$ is realized variance, N is the notional value and K_{VAR} is the price of variance. K_{VAR} is equal to the value of the portfolio that replicates the log contract. Demeterfi et al. (1999) derive formulae for valuing and pricing the variance swap, and directly obtaining the cost of the replicating portfolio. The key result, a formula for the fair value of future variance is given as Equation (26) of their paper, and shown here as Eq. (1).

$$K_{\text{var}} = \frac{2}{T}\left(rT - \left(\frac{S_0}{S_*}e^{rT} - 1\right) - \log\frac{S_*}{S_0} + e^{rT}\int_0^{S_*}\frac{1}{K^2}P(K)dK + e^{rT}\int_{S_*}^{\infty}\frac{1}{K^2}C(K)dK\right) \tag{1}$$

In Eq. (1) K_{var} is the fair price of future variance. S_0 is the underlying asset price at time 0. K is the strike price. $P(K)$ and $C(K)$ are the prices of out of the money calls and puts. S_* is the value of the underlying asset at the boundary between the calls and the puts (e.g. at the money). T is the maturity of the option.

Britten-Jones and Neuberger (2000) demonstrate how, given a continuum of European option prices with strikes and maturities ranging from zero to infinity, a condition can be derived which must be satisfied by all price processes consistent with the given set of option prices. Derman and Kani (1994), Dupire (1994, 1997), and Rubinstein (1994) showed that when volatility is deterministic, a unique price process exists that is consistent with option prices. Britten-Jones and Neuberger extended this analysis to a non-deterministic volatility setting, where many consistent price processes are possible. They derived their results in a discrete framework, using a time-price grid, and took limits as the interval sizes approach zero to obtain continuous counterparts. In an appendix, they also derived their results directly in a diffusion setting. Britten-Jones and Neuberger's simple condition is given as Equation (10) in Proposition 1 of their paper and is shown here as Eq. (1a).[1]

$$E\left[\left(\frac{S_{t+h} - S_t}{S_t}\right)^2 \mid S_t = K\right] = \frac{[C(t+h,K) - C(t,K)](u-1)^2(u+1)/u}{C(t,Ku) - (1+u)C(t,K) + uC(t,K/u)} \tag{1a}$$

In Eq. (1a) $C(t, K)$ is the call option price at strike price K and future time t, h is the size of the discrete time intervals used in the setting and u is the geometric factor acting on stock prices that determines the possible stock prices given the discrete time intervals.

The authors show that all price processes satisfying their Proposition 1 have the same (risk neutral) expectation for squared price volatility (e.g. price variance) over any given time period, and thus imply the same one-period forecast of volatility. Because this forecast is common to all such processes they refer to it as the "model-free" implied volatility. The analytical formula needed to extract the "model-free" implied volatility for a period between any two arbitrary future dates, from current prices of options expiring on those dates is given as Equation (13) of Proposition 2 of their paper shown here as Eq. (1b).[2] The authors note that this equation was derived independently by Carr and Madan (1998) in the context of pricing and hedging variance swaps, using results from Neuberger (1994), and the well-known Breeden and Litzenberger (1978) result.

$$E_0\left[\int_0^{t_2}\left(\frac{dS_t}{S_t}\right)^2\right] = 2\int_0^{\infty}\frac{C(t_2,K) - C(t_1 - K)}{K^2}dK \tag{1b}$$

A simplification of Eq. (1b) for the period between the current time and any arbitrary future date, is given as Equation (14) of Britten-Jones and Neuberger (2000), and is reproduced here as Eq. (2).

$$E_0\left[\int_0^{t_2}\left(\frac{dS_t}{S_t}\right)^2\right] = 2\int_0^{\infty}\frac{C(t_2,K) - \max(S_0 - K,0)}{K^2}dK \tag{2}$$

In Eq. (2) $\max(S_0 - K,0)$ is the intrinsic value of the option at time 0.

Bakshi and Madan (2000) observe that though the payoff functions of other securities are spanned by options, this has not resulted in a simplification of security valuations, because options themselves are complex to value. They propose the use of an alternative spanning entity. Namely; the characteristic function of the state price density (SPD), which they argue, significantly simplifies option pricing. The SPD is the discounted risk-neutral density function, and its characteristic function can be obtained via a Fourier transform.[3] Theorem 1 of Bakshi and Madan (2000) demonstrates that in an arbitrage free setting, the continuum of characteristic functions and the continuum of options are equivalent classes of spanning securities. It follows as a special case of Theorem 1 that all twice differentiable payoff functions can be algebraically spanned by a continuum of out-of-the-money calls and puts.

[1] Proposition 1: In any continuous risk-neutral process, the expectation of squared return, conditional on the stock price and time, is determined by the initial option prices as

$$E\left[\left(\frac{S_{t+h} - S_t}{S_t}\right)^2 \mid S_t = K\right] = \frac{[C(t+h,K) - C(t,K)](u-1)^2(u+1)/u}{C(t,Ku) - (1+u)C(t,K) + uC(t,K/u)} \tag{10}$$

The converse is also true; any continuous martingale process for S that satisfies the above condition for all $K \in \boldsymbol{K}$ and $t \in \boldsymbol{T}$ will price all European options correctly by their expected payoffs.

[2] Proposition 2: The risk-neutral expected sum of squared returns between two arbitrary dates t_1 and t_2 is given from the set of prices of options expiring on these two dates as

$$E_0\left[\int_0^{t_2}\left(\frac{dS_t}{S_t}\right)^2\right] = 2\int_0^{\infty}\frac{C(t_2,K) - C(t_1 - K)}{K^2}dK \tag{13}$$

[3] Recent research discusses how the use of wavelets has advantages over Fourier transforms in option pricing (Ortiz-Gracia & Oosterlee, 2013).