Factor Models for Option Pricing

Peter Carr, Principal

Banc of America Securities

Dilip Madan, Professor

University of Maryland

Abstract

Options on stocks are priced using information on index options and viewing stocks in a factor model as indirectly holding index risk. The method is particularly suited to developing quotations on stock options when these markets are relatively illiquid and one has a liquid index options market to judge the index risk. The pricing strategy is illustrated on IBM and Sony options viewed as holding SPX and Nikkei risk respectively.

Current Version: May 21, 2001
File Reference: vgsiapfm1pres.tex

Introduction and Overview

- Consider the problem of pricing a stock option when there are very few (possibly none) liquid options written on the same underlying.
- Suppose also that we have a liquid market for options on a major financial index which may or may not include the stock in question.
- We show how one can price options on the individual stocks by viewing the risk of the stock as the sum of the index risk and an idiosyncratic firm-specific component.
- The index risk is priced consistent with the index options market, while the idiosyncratic component is priced consistent with the true statistical probability, with no change of probability on this component.

Illustration of Approach

- Our methods requires the user to pick a model and a method for pricing index options.
- The model we illustrate our approach with is the variance gamma (VG) model of Madan, Carr, and Chang (1998).
- The method used is the Fast Fourier Transform (FFT) approach of Carr and Madan (1999).
- We illustrate factor option pricing by first pricing IBM options as indirect exposure to the S&P 500 index. Here, we may compare the results of factor option pricing with a direct calibration of the liquid options market in IBM options.
- Next, we price Sony options as indirect exposure to the Nikkei. Since Sony options are relatively illiquid, a comparison of the indirect calibration with the direct calibration is not possible.

Assumptions on Index Process

- We suppose that the statistical process for the index level at time t, I_t, is given by:

$$I_t = I_0 \exp\left((\mu - q)t + X_t + \omega_p t\right),$$

where μ is the assumed constant expected rate of return, q is the assumed constant dividend yield on the index, $X(t)$ is a purely discontinuous Lévy process, and ω_p compensates the exponential Lévy process, and is given by:

$$\omega_p = -\int_{-\infty}^{\infty}(e^x - 1)k_p(x)dx,$$

where k_p is the statistical Lévy density for the process $X(t)$.

- We also suppose that the risk-neutral process for the index level is given by:

$$I_t = I_0 \exp\left((r - q)t + X_t + \omega t\right),$$

where r is the assumed constant riskfree rate, and ω_q is given by:

$$\omega_q = -\int_{-\infty}^{\infty}(e^x - 1)k_q(x)dx,$$

where k_q is the risk-neutral Lévy density for the process $X(t)$.

Some Choices for the Index Driver

- Recall that the statistical and risk-neutral processes are both modelled as:

$$I_t = I_0 \exp\left((\mu - q)t + X_t + \omega t\right),$$

- Six possibilities for X_t are:
  1. The 2 parameter Finite Moment Log Stable (FMLS) process of Carr and Wu (2000)
  2. The 2 parameter symmetric VG process of Madan and Seneta (1980)
  3. Its 3 parameter generalization to the asymmetric VG model of Carr, Madan, and Chang (1998),
  4. The 4 parameter generalization of the asymmetric VG model to the CGMY model of Carr, Geman, Madan, and Yor (2000).
  5. The use of an inverse Gaussian time change and the normal inverse Gaussian model of Barndorff-Nielsen
  6. Its generalizations to the class of generalized hyperbolic distributions.
- For all 6 models, the characteristic function of X_t is known in closed form and thus option prices can be efficiently computed using our FFT approach.
- All 6 models capture fat tails and all but the second allow for non-zero skewness in index returns.
- All 6 models also provide high quality calibrations of the options data across the strike dimension.
- The first model is also designed to work across maturities.

- It appears that an analog of stochastic volatility (called stochastic arrival rate) is needed to make the latter 5 models work across the maturity dimension.

- For concreteness, we henceforth focus on the asymmetric VG model where the index driver X_t is described by 3 parameters, σ, ν, and θ:

$$X_t = \theta G_t(\nu) + \sigma W_{G_t(\nu)}.$$

- This index driver results from time-changing a Brownian motion with drift θ and variance rate σ^2 using a gamma process G_t with a mean rate of one and variance rate ν. In this case, the vector $\Theta = (\sigma, \nu, \theta)$.

- For a symmetric VG process, $\theta = 0$.

Individual Stock Statistical Process

- We suppose that continuously compounded returns on the stock are linearly related to index returns up to a zero mean idiosyncratic residual.
- To allow for fat tails in the residuals, we model the residual process as a symmetric VG process.
- Letting S_t be the stock price at time t, the statistical return process is assumed to be:

$$\ln\left(S_{t+h}/S_t\right) = \alpha + \beta \ln\left(I_{t+h}/I_t\right) + Y_{t+h}(s,\kappa) - Y_t(s,\kappa),$$

  where Y_t is an independent symmetric VG process given by time-changing a Brownian motion with variance rate s^2 by a gamma process with unit mean rate and variance rate κ.
- If we let W_t denote the standard Brownian motion, then:

$$Y_t(s,\kappa) = sW_{g_t(\kappa)},$$

  where $g_t(\kappa)$ is the gamma process with variance rate parameter κ.
- The above stock return model implies that the stock price process is:

$$S_t = S_0 \exp\left(\alpha t + \beta \ln(I_t/I_0) + Y_t\right),$$

  where we have dropped the dependence of Y on the parameters for notational convenience.

Individual Stock Risk-Neutral Process

- For pricing standard options on the stock, we must identify the risk-neutral density for the stock price.
- For this purpose, we suppose that the risk-neutral process for the index may be identified using the prices of index options in a liquid market for such options.
- For the firm-specific risk, if we have some options trading on the stock, then we may allow for a measure change in Y_t, but we suppose that the risk-neutral process is still in the VG class.
- Denoting the risk-neutral parameters of the process Y_t by σ_f, ν_f, and θ_f, the risk-neutral process for the stock is assumed to be:

$$S_t = S_0 \exp\left((r - \delta)t + \beta X_t(\sigma, \nu, \theta) + Y_t(\sigma_f, \nu_f, \theta_f) + \lambda t\right),$$

  where δ is the stock dividend yield and λ is the compensator for the jump processes in the exponential, and is determined explicitly below.
- We note that if the firm-specific risk is not priced, then the risk-neutral process for Y_t should be the same as the statistical process, and thus $\theta_f = 0, \sigma_f = s, \nu_f = \kappa$.
- In the absence of individual stock option prices, or conditional on supposing that idiosyncratic risk is not priced, we may infer the regression coefficient β by standard time series methods and set $(\sigma_f, \nu_f, \theta_f)$ to their statistical counterparts $(s, \kappa, 0)$, as identified by a statistical analysis.
- Alternatively, if we have some limited data on stock options, we may infer these parameters from such information.

FFT Index Option Pricing

- European-style index options may be easily priced using our FFT methodology, which is completely general.
- Our Fourier methodology analytically relates the generalized Fourier transform of an option price, considered as a function of log strike, to the risk-neutral characteristic function of the log index level at maturity.
- Hence, if the given risk-neutral characteristic function is analytic, then so is the generalized Fourier transform of the option price.
- The index option prices at all strike levels may then be obtained by a single Fourier inversion, possibly using FFT to speed up the computation.
- By vectorizing the above operation across maturities, the output is the desired matrix of option values across all strikes and maturities.
- In the VG case, we may suppose that this index option price is given by the pricing function:

$$w = W(I_0, K, r, q, \sigma, \nu, \theta, T, u), \qquad K > 0, T > 0,$$

where u is a dummy variable indicating whether the option is a call ($u = 1$ if a call and $u = 0$ if a put).

FFT Individual Stock Option Pricing

- The FFT methodology can also be used for (European) single name options, provided one has an analytic expression for the risk-neutral characteristic function of the log price at maturity.
- For our model, the independence of the processes X and Y implies that:

$$\begin{aligned} Ee^{iu\ln S_t} &= \exp\left(iu(\ln(S_0) + (r - \delta + \lambda)t)\right) X \\ &\quad \left(1 - i\theta\nu\beta u + \sigma^2\beta^2 u^2\nu/2\right)^{-t/\nu} X \\ &\quad \left(1 - i\theta_f\nu_f u + \sigma_f^2\nu_f u^2/2\right)^{-t/\nu_f}. \end{aligned}$$

- The coefficient λ is determined on noting that the risk-neutral characteristic function $Ee^{iu\ln S_t}$ evaluated at $u = -i$, must be $e^{(r-\delta)t}$.
- It follows that:

$$\lambda = \frac{1}{\nu}\ln\left(1 - \theta\nu\beta - \sigma^2\beta^2\nu/2\right) + \frac{1}{\nu_f}\ln\left(1 - i\theta_f\nu_f - \sigma_f^2\nu_f/2\right).$$

- We christen the pricing methodology described above as VGSI for Variance Gamma Stock Index.
- The matlab program $vgsi.m$, included in the appendix, uses the above characteristic function to compute individual stock option prices via the FFT methodology.
- A typical call to the pricing function is given by:

$$w = vgsi(S_0, K, r, \delta, \sigma, \nu, \theta, \beta, \sigma_f, \nu_f, \theta_f, T, u), \qquad K > 0, T > 0,$$

where u is one for a call and zero for a put.

Applying VG to Price IBM Options

- IBM options are liquid and so we can compute VG model prices. These can be used to test the error of the VGSI model prices, where the (very liquid) SPX index is used as our index.
- The data employed for this experiment were for July 1, 1998. The IBM stock price was at 116.8125 and we consider options of maturity .296 years. The risk-free rate for this maturity was 5.75 percent and the dividend yield was estimated at .63 percent. For this maturity, we had available 10 IBM put prices for strikes ranging from 95 dollars to 140 dollars and 9 IBM call prices with strikes ranging from 100 to 140 dollars.
- The VG model was estimated on this stock option data using only OTM prices, yielding parameter estimates of $.296, .1229, -.2882$ for σ, ν, θ respectively.

VGSI Pricing of IBM Options off SPX

- In applying VGSI, we first estimate the time series model of returns.
- For this purpose, we regressed daily log returns for IBM on SPX log returns over the period Dec. 31, 1993 to July 1, 1998, obtaining:

$$\ln(S_{t+1}/S_t) = .003 + 1.1920 \ln(I_{t+1}/I_t) + \varepsilon_{t+1}.$$

- The VG parameters for the residual may be estimated by the method of moments, with an estimate for s being the standard deviation of the residuals, and an estimate of κ given by the ratio of the fourth moment to three times s^4 less 1. The resulting annualized estimates for s and κ were .2688 and .0124 respectively (assuming 250 days in a year).
- The next step in the implementation of VGSI is the estimation of the risk-neutral density for the index risk. For this step, we obtained data on 25 SPX index options, with 6 calls and 19 puts at a maturity of .219 years. The risk-free rate was 5.75 percent and the dividend yield on the index was estimated at 1.45 percent. The level of the index at this date was $1,145.5643$. The risk-neutral VG parameter estimates for the SPX on this day for this maturity were $\sigma = .1704$, $\nu = .2363$, and $\theta = -.2143$.

IBM off SPX (con'd)

- If we now suppose that the risk-neutral estimate for the IBM beta is the same as the statistical estimate of 1.1920 and that the residual risk in the stock is not priced, then we have all the information needed to construct VGSI option prices for IBM.
- The results for these option prices are presented for puts in Figure 1 and for calls in Figure 2.
- We observe that the VG models fit is very good.
- Although the VGSI price curve follows the general shape of the other two curves, it overprices both the puts and the calls.
- We note that for pure jump processes like the VG, the statistical and risk-neutral volatilities and hence β estimates need not be the same. We advocate using the at-the-money option price to estimate the risk-neutral β. The dotted line in Figures 1 and 2 is the beta-adjusted VGSI price curve.
- We observe that we can effectively explain IBM option prices as indirectly pricing index risk with risk-neutral pricing of the residual risk, provided that we adjust risk-neutral betas to fit the at-the-money option price. Thus, VGSI is a potentially feasible methodology for pricing options indirectly as exposure to index factor risk.

VGSI Pricing of Sony Options off Nikkei

- Here we use the market prices of Nikkei index options, as well as the few option prices we have on Sony stock, to calibrate the risk-neutral parameters of the firm-specific process.

- The data used was from the market close of July 27, 1998. For the September 10 maturity, we have prices on three Sony options. In particular, we have a put struck at $12,000$ priced at 235, and two calls struck at $13,000$ and $14,000$ priced at 640 and 275 respectively. For these options, the time to expiration is .1205 years. The spot market price of Sony on this day was $13,010$.

- For the Nikkei index at a time to maturity of .126 years, with the spot at $15,908$ yen, the risk-free rate at .0062 per cent, and a dividend yield of .0003 per cent, we have prices on 4 index puts and 3 index calls. The puts were struck at $14,000, 14,500, 15,000$, and $15,500$ and were priced at $100, 160, 260$, and 410 respectively. The calls were struck at $16,000, 16,500$, and $17,000$ with prices $565, 335$, and 200 respectively.

- The VG parameter estimates were $\sigma = .2658, \nu = .0505$, and $\theta = -.3855$.

- Figure 3 presents a graph of the actual out-of-the-money option prices and the VG fitted curve. The implied level of volatility, skewness, and kurtosis in the risk-neutral density is respectively .2795, $-.2022$, and 3.1792.

- This information is also presented in the form of implied volatilities in Figure 4.

Sony off Nikkei (con'd)

- The VGSI model has 4 additional parameters associated with the residual risk. These are the stock beta, residual volatility, skewness, and kurtosis.
- The statistical beta was estimated from time series at .7593, while the residual volatility is .2261, with a kurtosis of .0045. We supposed a statistical skewness of zero in these calculations.
- For the corresponding risk-neutral values, we suppose that the residual risk may be treated risk-neutrally, and let the risk-neutral residual skew also be zero.
- However, we estimate the risk-neutral beta, volatility, and kurtosis using the three Sony options which we have available for this maturity. The estimated risk-neutral values were $\beta = .1931$, with a residual risk-neutral volatility of .3031,and an excess kurtosis parameter of 0.001.
- Figures 5 presents the actual Sony prices and the price curve as estimated by the VGSI option pricing model. The same information is presented in the form of Black Scholes implied volatilities in Figure 6.

Conclusion

- We presented a methodology for pricing options on stocks using a factor model, whereby stock risk is viewed as the sum of indirect index risk and residual risk.
- The method is particularly suited to developing quotations on stock options when these markets are relatively illiquid and one has a liquid index options market to judge the index risk.
- The method exploits the time series correlation between the stock and the index and also incorporates a statistical analysis of the higher moments of the regression residuals in the absence of data on individual stock options.
- Regression coefficients and parameters for the risk-neutral residual risk may be estimated if there is some information on a few stock options.
- The results provide market consistent quotes on the remaining range of strikes.
- Generally, we advocate the use of at-the-money options prices to calibrate the choice of the level for the risk-neutral beta coefficient.
- The pricing strategy is illustrated on IBM options without using the IBM option prices directly. We observe that the strategy of pricing options as an indirect exposure to index risk performs reasonably well.
- We also illustrate and report results for pricing Sony stock options off the Nikkei index.

Bibliography

- Barndorff-Nielsen, O.E. (1998), "Processes of normal inverse Gaussian type," *Finance and Stochastics*, 2, 41-68.
- Carr, P., H. Geman, and D.B. Madan (2001), "Pricing and Hedging in Incomplete Markets," forthcoming in the *Journal of Financial Economics.*
- Carr, P., H. Geman, D.B. Madan, and M. Yor (1999), "The Fine Structure of Asset Returns: An Empirical Investigation," forthcoming in *The Journal of Business.*
- Carr, P. and D. B. Madan (1998), "Option Valuation using the fast Fourier transform," *Journal of Computational Finance,* 2, 61-73.
- Eberlein, E. and U. Keller (1995), "Hyperbolic Distributions in Finance," *Bernoulli 1*, 3, 281-299.
- Eberlein, E., U. Keller, and K. Prause (1998), "New Insights into Smile, Mispricing and Value at Risk: The Hyperbolic Model," *Journal of Business*, 71, 371-406.
- Madan, D.B., Carr, P., and E. Chang (1998), "The Variance Gamma Process and Option Pricing Model," *European Finance Review*, 2, 79-105.

Appendix

Matlab Programs for computing VGSI prices. The procedure computes the inverse Fourier transform of the Fourier transform of the modified call price to obtain call prices which are then transformed by put-call parity for put prices if desired. The inputs are

1. p, the current spot price.
2. x, the option strike.
3. r, the interest rate.
4. q, the dividend yield.
5. sig, nu, th the index VG parameters.
6. bta, the asset beta with respect to the index.
7. sgr, nur, thr, the VG parameters for the idiosyncratic residual.
8. t, the option maturity
9. u, a flag that is 1 for calls and 0 for puts.

The output is the option price. The program calls vgsifr.m, that calls in turn vgsifrf.m and vgsicf.m. The final program is the characteristic function for the log of stock price at maturity.

```
function y = vgsi(p,x,r,q,sig,nu,th,bta,sgr,nur,thr,t,u)

u1=vgsifr(log(x),p,r,q,t,sig,nu,th,bta,sgr,nur,thr);

uu=u1+(1-u).*(x.*exp(-r.*t)-p.*exp(-q.*t)); y=uu;

return

function y=vgsifr(k,p,r,q,t,sig,nu,th,bta,sgr,nur,thr)

aa=[k p r q t];

rni=[sig nu th];

sip=[bta sgr nur thr];

nn1=8; nn2=64; nn3=256;

uu1=quad8('vgsifrf',0,nn1,[],[],aa,rni,sip);

uu2=quad8('vgsifrf',nn1,nn2,[],[],aa,rni,sip);

uu3=quad8('vgsifrf',nn2,nn3,[],[],aa,rni,sip);

uu=uu1+uu2+uu3;

y=2.*uu;

return
```

```
function y = vgsifrf(v,P1,P2,P3)

k=P1(1);

p=P1(2);

r=P1(3);

q=P1(4);

t=P1(5);

sig=P2(1);

nu=P2(2);

th=P2(3);

bta=P3(1);

sgr=P3(2);

nur=P3(3);

thr=P3(4);

aa=.03;

u1=exp(-r.*t);

u2=vgsicf(v-(aa+1).*i,p,r,q,t,sig,nu,th,bta,sgr,nur,thr);

u3=aa.^2+aa-v.^2+i.*v.*(2.*aa+1);

u4=exp(-aa.*k);

u5=exp(-i.*v.*k);

u6=1./(2.*pi);
```

```
uu=u4.*u6.*u5.*u1.*u2./u3;

y=real(uu);

return

function y = vgsicf(u,p,r,q,t,sig,nu,th,bta,sgr,nur,thr)

h1=(1./nu).*log(1-bta.*th.*nu-sig.^2.*bta.^2.*nu./2);

h2=(1./nur).*log(1-thr.*nur-sgr.^2.*nur./2);

lda=h1+h2;

u1=log(p)+(r-q).*t+lda.*t;

u2=exp(i.*u.*u1);

u3=1-i.*th.*nu.*bta.*u+(sig.^2.*bta.^2.*nu./2).*u.^2;

u4=u3.^(-t./nu);

u5=1-i.*thr.*nur.*u+(sgr.^2.*nur./2).*u.^2;

u6=u5.^(-t./nur);

uu=u2.*u4.*u6;

y=uu;

return
```