Peter Carr
New York University

Hélyette Geman
Université Paris Dauphine and Ecole Superieure des Sciences Economiques et Commerciales

Dilip B. Madan
University of Maryland

Marc Yor
Université Paris 6.

The Fine Structure of Asset Returns: An Empirical Investigation*

We investigate the importance of diffusion and jumps in a new model for asset returns. In contrast to standard models, we allow for jump components displaying finite or infinite activity and variation. Empirical investigations of time series indicate that index dynamics are devoid of a diffusion component, which may be present in the dynamics of individual stocks. This leads to the conjecture, confirmed on options data, that the risk-neutral process should be free of a diffusion component. We conclude that the statistical and risk-neutral processes for equity prices are pure jump processes of infinite activity and finite variation.

I. Introduction

Asset returns have been modeled in continuous time as diffusions by Black and Scholes (1973) and Merton (1973), as pure jump processes by Cox and Ross (1976), and as jump-diffusions by Merton (1976). The jump processes studied by Cox and Ross display finite activity, while some recent research has considered some pure jump processes with infinite activity. Two examples of these infinite-activity pure jump processes

* We would like to thank seminar participants at Princeton University, the University of Aarhus, University of Freiburg, Eidengenössische Technische Hochschule Zurich, the University of Chicago, the University of Massachusetts—Amherst, the International Center for Business Information (ICBI) Global Derivatives 2000 meeting, the meeting of the World Congress of Mathematicians in Barcelona, the workshop in mathematical finance at Carnegie Mellon, and the conference in financial markets at the University of Warwick, for their comments and discussion, and, in particular, David Heath, Ajay Khanna, Sebastian Raible, and Liuren Wu. Dilip Madan would also like to acknowledge Centre de Recherches sur la Gestion at the Université Paris IX-Dauphine for their support of his visit, where he completed part of the work on this project.

(*Journal of Business,* 2002, vol. 75, no. 2)

0021-9398/2002/7502-0004$10.00

are the variance gamma model studied by Madan and Seneta (1990) and Madan, Carr, and Chang (1998) and the hyperbolic model considered in Eberlein, Keller, and Prause (1998). The rationale usually given for describing asset returns as jump-diffusions is that diffusions capture frequent small moves, while jumps capture rare large moves. Given the ability of infinite-activity jump processes to capture both frequent small moves and rare large moves, the question arises as to whether it is necessary to employ a diffusion component when modeling asset returns.

To answer this question, this article develops a continuous time model that allows for both diffusions and for jumps of both finite and infinite activity. We define a pure jump process to be one of finite/infinite activity if the number of price jumps in any interval of time is finite/infinite. The parameters of our process further allow the jump component to have either finite or infinite variation. Thus, our model synthesizes the features of the above-cited continuous-time models and captures their essential differences in parametric special cases. The model is called the CGMY model, after the authors of this article. We employ this model to study both the statistical process needed to assess risk and allocate investments and the risk-neutral process used for pricing and hedging derivatives. Our process generates a closed-form expression for the characteristic function of log prices but not for the return density. Nonetheless, by employing our methodology on a time series of stock returns and options data, we demonstrate how knowledge of the characteristic function can be used to infer econometrically the fine structure of the statistical and risk-neutral processes.

We find that index returns tend to be pure jump processes of infinite activity and finite variation, both statistically and risk-neutrally. Thus, the index return processes appear to have effectively diversified away any diffusion risk that may be present in individual stock returns. We note, however, that even the diffusion components estimated in individual equity returns appear to be statistically insignificant. In contrast, the jump components account for consistently significant skewness levels that statistically may be either positive or negative but that risk-neutrally are consistently negative. The signature profile for the mean corrected density for asset returns appears to be a long spike near zero, conjoined with two convex fans describing larger returns. The departure in shape from the Gaussian is quite glaring, as the normal distribution is always concave within one standard deviation of the mean. In contrast, the densities of high-activity finite variation processes are consistent with the data from both time series and option prices. We also note that, since dynamic trading often results in profit and loss distributions similar to those generated by our process, our research should be relevant to the literature on prescribing capital requirements and on designing insurance contracts covering hedge fund losses.

Thus, the contribution of our article is threefold. First, on the theoretical side, we introduce a new stochastic process, which we use to describe asset returns and model option values. Second, on the computational side, we demonstrate the use of Fourier inversion via the Fast Fourier transform as a technique for determining statistical and risk-neutral densities numerically. Finally, on the empirical side, we show that one can usually dispense with diffusions in describing the fine structure of asset returns, so long as the jump process used is one of infinite activity and finite variation.

From our estimates of the statistical and risk-neutral processes for each of a set of names and indices, we also offer some preliminary conjectures on the suggested nature of the implied measure change. A definite conclusion in this direction must await a systematic empirical investigation that jointly estimates the statistical and risk-neutral processes on the same data. To model the measure change adequately, parametric restrictions imposed by the requirements of equivalence for the two measures are undesirable. Thus, it is instructive to construct the measure change using approximating finite activity processes that truncate the very small and very large jumps. Our findings are informative as to the relevant theoretical directions such research may take. In our tentative view, a critical input for constructing the measure change is the structure of open interest in the options market. We hypothesize that large open interest in out-of-the-money puts are a possible source of the negative skewness observed in option-implied distributions.

We recognize that option pricing for processes with pure jump components forces a move out of the traditional realm of arbitrage pricing into the domain of equilibrium pricing. On the positive side, our setting allows us to use option prices to study the measure change and the nature of the underlying equilibrium. Furthermore, arbitrage pricing can still be used to value more complex claims relative to option prices, even though prices jump.

The outline of the article is as follows. In Section II, we present the details of the synthesizing model and its parametric properties. In Section III, we define the statistical and risk-neutral stock price model when the underlying uncertainty is a Lévy process. Section IV provides analytical details for constructing the higher moments, decomposing expected total variation into its diffusion and pure jump components, and explicitly illustrating the measure change process. In Section V, we present the estimation methodology and the results, and, in Section VI, we discuss the results from a variety of perspectives. Section VII concludes.

II. The CGMY Model

Before describing our model, we describe our mechanism for inferring continuous time sample path properties from discrete observations. We recognize that this inference is difficult and fraught with peril. After all, how is one to infer from daily observations, whether the price process has discontinuities and, if it does, how many? Our path to the finer structure of asset returns, measured by the log price relative, is through the characteristic function for the logarithm of the stock price. The Lévy Khintchine theorem uniquely represents this characteristic function for infinitely divisible processes. Armed with this fundamental result and some modern computational advances in Fourier inversion, maximum likelihood estimation of the parameters of the statistical process from time-series data becomes feasible. Furthermore, similar methods may be employed to estimate risk-neutral parameters from options data, as is shown in Carr and Madan (1998). We are thus able to design a probe of the data that enables one to learn about the fine structure of asset returns from discrete observations, admittedly under some maintained auxiliary hypotheses.

The starting point of our analysis is the geometric Brownian motion model of Black and Scholes (1973) and Merton (1973), in which the cumulative return is modeled as the Lévy process given by arithmetic Brownian motion. We seek to replace this process with one that enjoys all of the fundamental properties of Brownian motion, except for pathwise continuity and scaling, but that permits a richer array of variation in higher moment structure, especially at shorter horizons. These considerations lead us to focus on the auxiliary hypotheses embedded in infinitely divisible processes of independent and homogeneous increments. For reasons we will outline later, we are also interested in processes with finite-variation jump components. For such processes, the characteristic function is uniquely characterized by the Lévy Khintchine theorem in terms of the drift rate a, the diffusion coefficient b, and the Lévy density $k(x)$. Specifically, if $X(t)$ is an infinitely divisible process with a finite variation jump component and independent and homogeneous increments, then its characteristic function is uniquely given by

$$E\left[e^{iuX(t)}\right] = \exp\left[iuat - \frac{u^2 b^2 t}{2} + t\int_{-\infty}^{\infty} (e^{iux} - 1)k(x)\ dx\right].$$

Heuristically, the Lévy density measures by $k(x)\ dx$ the arrival rate of jumps of size x. The jump component of such processes is completely characterized by this Lévy density. Our modeling focus is on candidate parametric choices for this Lévy density, and so we begin our analysis by considering pure jump processes.

The next subsection presents the details of the variance gamma model developed by Madan and Seneta (1990) and extended to incorporate skewness by Madan and Milne (1991) and Madan et al. (1998). The Madan et al. paper shows that this model permits a parsimonious description of the volatility smile observed in option prices at all maturities and for a wide variety of underlying assets. The results of this study suggest that the success of the variance gamma process in explaining the smile is likely due to the fact that the process is a pure jump process, which displays infinite activity but finite variation. The following subsection develops the CGMY process that generalizes the variance gamma process by adding a parameter permitting finite or infinite activity and finite or infinite variation.

A. The Variance Gamma Process

There are two representations for the variance gamma process, both of which are useful, but in different contexts. In the first representation, which gave rise to the name, the variance gamma process is interpreted as a Brownian motion with drift, where time is changed by a gamma process. Let $W(t)$ be a standard Brownian motion, and let $G(t;\ 1, \nu)$ be an independent gamma process with mean rate unity and variance rate ν. The density of the gamma process at time t is given by

$$f(g) = \frac{g^{t/\nu - 1} \exp\left(-\frac{g}{\nu}\right)}{\nu^{t/\nu} \Gamma\left(\frac{t}{\nu}\right)}, \tag{1}$$

while the characteristic function is given by

$$\phi_G(u, t) = E\{\exp[iuG(t)]\} = \left(\frac{1}{1 - i\nu u}\right)^{t/\nu}. \tag{2}$$

The variance gamma process has three parameters, σ, ν, and θ, and the process $X_{VG}(t;\ \sigma, \nu, \theta)$ is given by

$$X_{VG}(t;\ \sigma, \nu, \theta) = \theta G(t;\ \nu) + \sigma W[G(t;\ \nu)]. \tag{3}$$

The variance gamma (VG) process has a particularly simple characteristic function:

$$\phi_{VG}(u, t) = E\{\exp[iuX_{VG}(t)]\} = \left(\frac{1}{1 - i\theta\nu u + \sigma^2 \nu u^2/2}\right)^{t/\nu}. \tag{4}$$

This characteristic function is easily obtained from (2) by conditioning on the gamma time and using the fact that the conditioned random variable is Gaussian.

For the second representation, the VG process is interpreted as the difference of two independent gamma processes, since the characteristic function factors, using the fact that

$$\frac{1}{1 - i\theta\nu u + \sigma^2 \nu u^2/2} = \left(\frac{1}{1 - i\eta_p u}\right)\left(\frac{1}{1 + i\eta_n u}\right),$$

where η_p, η_n satisfy

$$\eta_p - \eta_n = \theta\nu,$$

$$\eta_p \eta_n = \frac{\sigma^2 \nu}{2}.$$

It follows that η_p, $-\eta_n$ are the roots of the equation

$$x^2 - \theta\nu x - \sigma^2 \nu/2 = 0,$$

whereby

$$\eta_p = \sqrt{\frac{\theta^2 \nu^2}{4} + \frac{\sigma^2 \nu}{2}} + \frac{\theta \nu}{2},$$

$$\eta_n = \sqrt{\frac{\theta^2 \nu^2}{4} + \frac{\sigma^2 \nu}{2}} - \frac{\theta \nu}{2}.$$

The two gamma processes may be denoted $G_p(t;\ \mu_p, \nu_p)$ and $G_n(t;\ \mu_n, \nu_n)$, with, respectively, mean and variance rates μ_p, μ_n and ν_p, ν_n. For these gamma processes, we have that $\mu_p = \eta_p/\nu$, $\mu_n = \eta_n/\nu$, while $\nu_p = \mu_p^2 \nu$, and $\nu_n^2 = \mu_n^2 \nu$. We note that the ratio of the variance rate to the square of the mean rate is the same for both gamma processes and is equal to ν. We then have that

$$X_{VG}(t;\ \sigma, \nu, \theta) \overset{law}{=} G_p(t;\ \mu_p, \nu_p) - G_n(t;\ \mu_n, \nu_n). \tag{5}$$

From this representation of the VG process and classical representations for the Lévy measures of gamma processes, Madan et al. (1998) show that the Lévy density for the VG process is

$$k_{VG}(x) = \begin{cases} \dfrac{\mu_n^2}{\nu_n} \dfrac{\exp\left(-\frac{\mu_n}{\nu_n}|x|\right)}{|x|} & for\ x < 0 \\ \dfrac{\mu_p^2}{\nu_p} \dfrac{\exp\left(-\frac{\mu_p}{\nu_p}|x|\right)}{|x|} & for\ x > 0. \end{cases} \tag{6}$$

The division by the absolute value of the jump size in the VG Lévy density (6) results in a process of infinite activity, as the VG Lévy measure integrates to infinity. It is also clear that, since $|x|$ is integrable with respect to the VG Lévy density, the process is one of finite variation.

B. The CGMY Process

In this subsection, we generalize the VG Lévy density to the CGMY Lévy density with parameters C, G, M, and Y. Specifically, the Lévy density of the CGMY process $k_{\mathrm{CGMY}}(x)$ is given by

$$k_{\mathrm{CGMY}}(x) = \begin{cases} C\dfrac{\exp(-G|x|)}{|x|^{1+Y}} & \text{for } x < 0 \\ C\dfrac{\exp(-M|x|)}{|x|^{1+Y}}, & \text{for } x > 0, \end{cases} \tag{7}$$

where $C > 0$, $G \geq 0$, $M \geq 0$, and $Y < 2$. The condition $Y < 2$ is induced by the requirement that Lévy densities integrate x^2 in the neighborhood of 0. We denote by $X_{\mathrm{CGMY}}(t;\ C, G, M, Y)$ the infinitely divisible process of independent increments with Lévy density given by (7). The case $Y = 0$ is the special case of the VG process with the parameter identification

$$C = \frac{1}{\nu}, \tag{8}$$

$$G = \frac{1}{\eta_n}, \tag{9}$$

$$M = \frac{1}{\eta_p}. \tag{10}$$

These parameters play an important role in capturing various aspects of the stochastic process under study. The parameter C may be viewed as a measure of the overall level of activity. Keeping the other parameters constant and integrating over all moves exceeding a small level, we see that the aggregrate activity level may be calibrated through movements in C. For example, if one were to construct a model with a stochastic aggregate activity rate, then one could model C as an independent positive process, possibly following a square root law of its own. In the special case when $G = M$, the Lévy measure is symmetric, and, in this case, Madan et al. (1998) show that the parameter C provides control over the kurtosis of the distribution of $X(t)$. The case $G = M$ has also been studied by Koponen (1995), who gives an alternative expression for the characteristic function.

The parameters G and M, respectively, control the rate of exponential decay on the right and left of the Lévy density, leading to skewed distributions when they are unequal. For $G < M$, the left tail of the distribution for $X(t)$ is heavier than the right tail, which is consistent with the risk-neutral distribution typically implied from option prices. Thus, when G and M are implied from the risk-neutral distribution, their difference calibrates the price of a fall relative to a rise, while their sum measures the price of a large move relative to a small one. In contrast, in the statistical distribution, the difference between G and M determines the relative frequency of drops relative to rises, while their sum measures the frequency of large moves relative to small ones. The exponential factor in the numerator of the Lévy density leads to the finiteness of all moments for the process $X(t)$. As we typically construct a process at the return level, it is reasonable to enforce finiteness of the moments at this level.

The parameter Y was studied in Vershik and Yor (1995), and it arises in the process for the stable law. The parameter Y is particularly useful in characterizing the fine structure of the stochastic process. For example, one may ask whether the up jumps and down jumps of the process have a completely monotone Lévy density, and whether the process has finite or infinite activity, or variation. We briefly describe these properties.

Completely monotone Lévy density. A completely monotone (CM) Lévy density structurally relates arrival rates of large jump sizes to smaller jump sizes by requiring, among other things, that large jumps arrive less frequently than small jumps. Completely monotone Lévy densities are essentially mixtures of exponential functions by virtue of Bernstein's theorem, which shows that all such densities may be written in the form

$$k(x) = \int_0^\infty e^{-ax}\zeta(da) \tag{11}$$

for some positive measure ζ. In the sequel, we shall be concerned with measures that are absolutely continuous with respect to Lebesgue measure and $\zeta(da) = w(a)\, da$ for some positive weighting function $w(a)$. This restriction on Lévy densities is useful in limiting the class of pure jump models one may entertain, and the condition is intuitively a reasonable one. For a variety of other models along these lines, the reader is referred to Geman, Madan, and Yor (2001).

Finite variation process. From the perspective of option pricing theory, processes of finite variation (FV) or finite activity (FA) are potentially more useful in explaining the measure change from the statistical to the risk-neutral process, as they permit greater flexibility between the local characteristics of the martingale components under the two measures. For example, for infinite variation processes like Brownian motion, the volatility, and hence the local martingale component, is invariant under an equivalent change of measure. For infinite variation jump processes, like the stable laws with exponent above unity, equivalence of the measure change implies (see Jacod and Shiryaev 1980, condition 3.25, p. 160) that the difference between the risk-neutral and statistical Lévy densities be of finite variation, and this imposes the restriction that the two processes have the same exponent, or, heuristically speaking, that they be of infinite variation in the same way. Clearly, if the processes are themselves of finite variation, then the difference in the Lévy densities will also be of finite variation, and, hence, no parametric restriction is required on account of this condition. These observations are important in light of the evidence from time series and from options data that indicates that risk-neutral volatilities are substantially higher than their statistical counterparts.

The flexibility of FA processes is even greater than that of infinite-activity, finite-variation processes, since Jacod and Shiryaev (1980) (see condition 4.39, c, (v), p. 246) show that parametric restrictions may also be imposed by requiring equivalence for the latter class of processes. Equivalence essentially requires that the Hellinger distance between the Lévy densities be finite. In particular, one may not have one process be of finite activity while the other is of infinite activity. Heuristically, one may say that the two processes must be of infinite activity in the same way. For the specific case of CGMY, one may not change C or Y under an equivalent measure change.

If, however, the data suggest that these parameters do change, it is reasonable

TABLE 1 **Process Properties and Ranges for the Parameter *Y***

Range of Y Values	Properties of Process
$Y < -1$	Not completely monotone, finite activity
$-1 < Y < 0$	Completely monotone, finite activity
$0 < Y < 1$	Completely monotone, infinite activity, finite variation
$1 < Y < 2$	Completely monotone, infinite variation, finite quadratic variation

NOTE.—The parameter Y describes the behavior of the Lévy density near zero. In particular, for Y exceeding negative unity, Y is the rate at which this density tends to infinity near the origin.

to drop down to an approximating class of FA processes and to view the Lévy process models as truncated in a small neighborhood of zero. The required integrability conditions are then satisfied. From such a perspective, the measure change may always be constructed in the complement of a neighborhood of zero. The resulting advantage from an empirical standpoint is that one may freely calibrate all parameters to the respective statistical and risk-neutral data and then learn the nature of the measure change made by the market on the approximating finite activity process.

Finite activity process. Processes of FA are of interest as one may wish to group assets by their activity levels. Thus, the use of infinite activity processes in mathematical finance is best viewed as a first approximation designed to study highly liquid markets with large activity. The properties described above are all related to values for Y being in certain regions that are described in table 1. These properties will be demonstrated formally in theorem 2.

The motivation for pursuing closed-form solutions for densities and for option prices is frequently that the model then constitutes a tractable probe for data, permitting real-time parameter estimation from time-series returns and from option prices. Although we do not have closed forms for these entities in the case of the CGMY model, we are able to exploit the fact that the characteristic function of the process is available in closed form. Theorem 1 below displays the required characteristic function.

THEOREM 1. The characteristic function for the infinitely divisible process with independent increments and the CGYM Lévy density (7) is given by

$$\begin{aligned}\phi_{\mathrm{CGMY}}(u, t;\ C, G, M, Y) \\ = \exp\{tC\Gamma(-Y)[(M - iu)^Y - M^Y + (G + iu)^Y - G^Y]\}.\end{aligned}$$

Proof. From the Lévy Khintchine theorem, we have that

$$\phi_{\mathrm{CGMY}}(u, t) = \exp\left[t\int_{-\infty}^{\infty} (e^{iux} - 1)k_{\mathrm{CGMY}}(x)\ dx\right].$$

The integral in the exponent may be written as the sum of two integrals of the form

$$\int_0^\infty (e^{iux} - 1) = C\frac{\exp(-\beta x)}{x^{1+Y}}\ dx$$

for β equal to G and M, respectively, with iu replaced by $-iu$ for $\beta = G$. This integration may be performed as follows:

$$\begin{aligned}\int_0^\infty & Cx^{-Y-1}\{\exp[-(\beta - iu)x] - \exp(-\beta x)\}\ dx \\ &= C\int_0^\infty (\beta - iu)^Y w^{-Y-1}\exp(-w)\ dw \\ &\quad - C\int_0^\infty \beta^Y w^{-Y-1}\exp(-w)\ dw \\ &= C\Gamma(-Y)[(\beta - iu)^Y - \beta^Y].\end{aligned}$$

The result follows on substituting M and G for β and evaluating the case $\beta = G$ at $-iu$. Q.E.D.

THEOREM 2. The CGMY process

i) has a completely monotone Lévy density for $Y > -1$;
ii) is a process of infinite activity for $Y > 0$; and
iii) is a process of infinite variation for $Y > 1$.

Proof. For property i, we note that, for $Y < -1$, the quantity $1 + Y$ is negative and the Lévy density $x^{-(l+Y)}\exp(-\beta x)$ for $\beta = G, M$ increases near zero and then declines to zero as x tends to infinity. Hence, the density is clearly not completely monotone. When $(1 + Y) > 0$, we may write

$$\frac{1}{x^{1+Y}}\exp(-\beta x) = \int_\beta^\infty \frac{(a-\beta)^Y}{\Gamma(1+Y)}e^{-ax}\ da,$$

whereby we have complete monotonicity with weighting function $1_{a>\beta}(a - \beta)^Y/\Gamma(1 + Y)$.

For property ii, we note that, for negative values of Y, the Lévy measure integrates to a finite value in the neighborhood of zero, and so we have a process of finite activity. When Y exceeds zero, however, the Lévy measure integrates to infinity near zero and we have an infinite activity process.

For property iii, we note that $|x|k_{\text{CGMY}}(x)$ has a finite integral near zero for $Y < 1$, while this integral is infinite for $Y > 1$. Q.E.D.

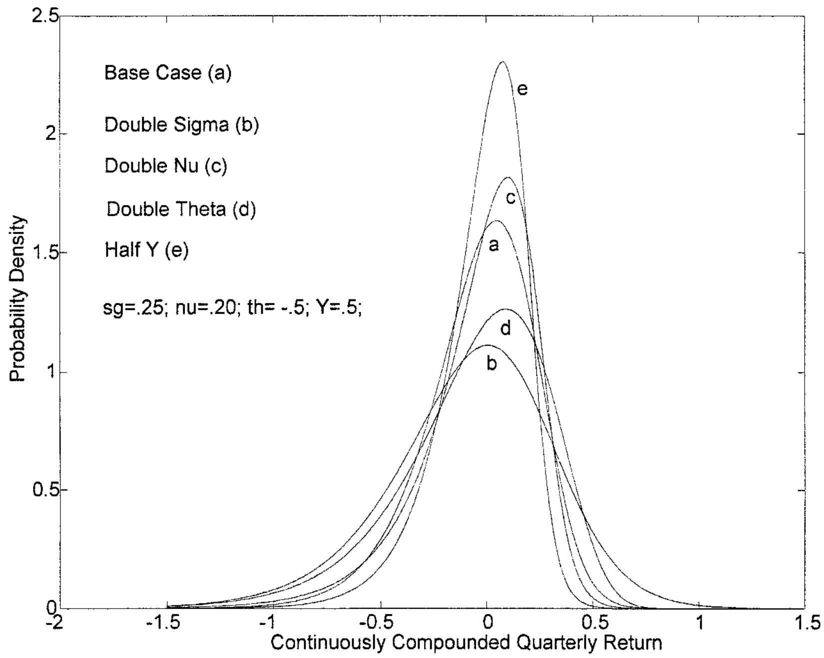

FIG. 1.—Densities of the CGMY model

III. The CGMY Stock Price Process

We model the martingale component of the logarithm of the stock price by the CGMY process. This is a fairly robust parametric class of stochastic processes consistent with a wide range of possible return distributions over finite holding periods. Besides being capable of calibrating to various levels of skewness and kurtosis, the CGMY model can also be used to study the nature of the fine structure of the stochastic process, as reflected in the parameter Y. To appreciate the breadth of possible densities, figure 1 graphs the density for log quarterly returns in various parameter settings. The parameters are referred to in their VG formulation with Y as the additional parameter. We present in curve a the base case for $\sigma = .25$, $\nu = .2$, and $\theta = -.5$, a typical setting for the S&P 500 index (henceforth SPX). We also initially set $Y = 0.5$. The other curves are double σ (curve b), double ν (curve c), double θ (curve d), and half Y (curve e). A variety of possible shapes and departures from normality may be observed.

A. The Statistical Stock Price Process

The CGMY model assumes that the martingale component of the movement in the logarithm of prices is given by the CGMY process. Hence, the stock price dynamics are assumed to be given by

$$S(t) = S(0)\exp[(\mu + \omega)t + X_{\mathrm{CGYM}}(t;\ C, G, M, Y)], \tag{12}$$

where μ is the mean rate of return on the stock and ω is a "convexity correction," defined by

$$\exp(-\omega t) = \phi_{\mathrm{CGYM}}(-i, t;\ C, G, M, Y). \tag{13}$$

Equations (12) and (13) define the evolution of the statistical process for the stock price. With a view to assessing the relevance of an additional diffusion component in our context, we next extend the model to include an orthogonal diffusion component. Define the extended CGMY process as

$$X_{\mathrm{CGMY}e}(t;\ C, G, M, Y, \eta) = X_{\mathrm{CGYM}}(t;\ C, G, M, Y) + \eta W(t),$$

where $W(t)$ is a standard Brownian motion independent of the process $X_{\mathrm{CGMY}}(t;\ C, G, M, Y)$. The extended stock price process has statistical dynamics given by

$$S(t) = S(0)\exp[(\mu + \omega - \eta^2/2)t + X_{\mathrm{CGMY}e}(t;\ C, G, M, Y, \eta)]. \tag{14}$$

The characteristic function for the logarithm of the stock price in this diffusion extended CGMY model is given by

$$\begin{aligned}\phi_{\ln(S)}(u, t) = {} & \exp\big(iu\{\ln[S(0)] + (\mu + \omega - \eta^2/2)t\}\big)\\ & \times \phi_{\mathrm{CGMY}}(u;\ C, G, M, Y)\exp(-\eta^2 u^2/2).\end{aligned} \tag{15}$$

Our statistical analysis employs the characteristic function (15) for the analysis of the time series of stock returns.

B. The Risk-Neutral Stock Price Process

We assume that the risk-neutral process for the stock lies in the robust five-parameter class of the diffusion-extended CGMY model, with a mean risk-neutral return given by the interest rate. The risk-neutral parameters can differ from their statistical counterparts, and, hence, they are denoted by $\tilde{C}$, $\tilde{G}$, $\tilde{M}$, $\tilde{Y}$, and $\tilde{\eta}$. Letting r denote the continuously compounded interest rate, the risk-neutral stock price process is

$$S(t) = S(0)\exp[(r + \tilde{\omega} - \tilde{\eta}^2/2)t + X_{\mathrm{CGMY}e}(t;\ \tilde{C}, \tilde{G}, \tilde{M}, \tilde{Y}, \tilde{\eta})], \tag{16}$$

with the characteristic function for the log of the stock price at time t given by

$$\begin{aligned}\tilde{\phi}_{\ln(S)}(u, t) = {} & \exp\big(iu\{\ln[S(0)] + (r + \tilde{\omega} - \tilde{\eta}^2/2)t\}\big)\\ & \times \phi_{\mathrm{CGMY}}(u;\ \tilde{C}, \tilde{G}, \tilde{M}, \tilde{Y})\exp(-\tilde{\eta}^2 u^2/2),\end{aligned} \tag{17}$$

and $\tilde{\omega}$ defined by

$$\exp(-\tilde{\omega}t) = \phi_{\mathrm{CGMY}}(-i, t;\ \tilde{C}, \tilde{G}, \tilde{M}, \tilde{Y}).$$

The parameters $\tilde{C}$, $\tilde{G}$, $\tilde{M}$, $\tilde{Y}$, and $\tilde{\eta}$ are the corresponding risk-neutral parameters estimated using data on option prices.

IV. Higher Moments, Total Variation, and Measure Changes

Once the statistical and risk-neutral processes have been estimated, we will have estimates for the parameters C, G, M, Y, and η, and for their risk-neutral equivalents $\tilde{C}$, $\tilde{G}$, $\tilde{M}$, $\tilde{Y}$, and $\tilde{\eta}$. Armed with these parameter estimates, we can determine the skewness and kurtosis under both the statistical and risk-neutral densities. We are also interested in assessing the relative magnitudes of the jump and diffusion components. We propose to measure this relative magnitude on the basis of the proportion of total quadratic variation contributed by each component. The quadratic variation of the diffusion component is clear, and we determine here the quadratic variation of the general CGMY component. We are also interested in the process for the measure change and wish to illustrate this process explicitly. For the higher moments, we develop explicit formulas for these in terms of the parameters.

A. Higher Moments of the $CGMY_e$ Process

The higher moments of the process may be obtained on successive differentiation of the characteristic function. For a general Lévy density $k(x)$ and diffusion coefficient η, one may show by differentiation that, for the random variable X representing the level of a Lévy process at time 1, we have

$$E\{[X - E(X)]^2\} = \eta^2 + \int_{-\infty}^{\infty} x^2 k(x)\, dx,$$

$$E\{[X - E(X)]^3\} = \eta^2 + \int_{-\infty}^{\infty} x^3 k(x)\, dx,$$

$$E\{[X - E(X)]^4\} = 3\big(E\{[X - E(X)]^2\}\big)^2 + \int_{-\infty}^{\infty} x^4 k(x)\, dx.$$

It follows, for the CGMY_e in particular, that

$$\text{Variance} = \eta^2 + C\Gamma(2-Y)\left(\frac{1}{M^{2-T}} + \frac{1}{G^{2-Y}}\right), \tag{18}$$

$$\text{Skewness} = \frac{C\Gamma(3-Y)\left(\frac{1}{M^{3-Y}} + \frac{1}{G^{3-Y}}\right)}{(\text{Variance})^{3/2}}, \tag{19}$$

$$\text{Kurtosis} = 3 + \frac{C\Gamma(4-Y)\left(\frac{1}{M^{4-Y}} + \frac{1}{G^{4-Y}}\right)}{(\text{Variance})^{2}}. \tag{20}$$

B. Decomposition of Quadratic Variation

We focus attention on the statistical process, with similar calculations applying to the risk-neutral case. The total quadratic variation over the interval $(0, t)$ of the diffusion component in the extended CGMY model with characteristic function (15) is $\eta^2 t$. For the jump component, the total quadratic variation is random, but its predictable quadratic variation and expectation is given by

$$t\int_0^\infty x^2 C\frac{\exp(-Mx)}{x^{1+Y}}\,dx + t\int_0^\infty x^2 C\frac{\exp(-Gx)}{x^{1+Y}}\,dx$$

$$= tC\Gamma(2-Y)\left(\frac{1}{M^{2-Y}} + \frac{1}{G^{2-Y}}\right). \tag{21}$$

We shall use equation (21) in computing the decomposition of quadratic variation reported later in our empirical results.

C. Measure Changes

The process for the Radon-Nikodym derivative of one measure with respect to another is not very interesting or informative when the underlying filtration is a diffusion with no jump component. On the other hand, for pure jump processes, Jacod and Shiryaev (1980) show how the change of measure process can be explicitly computed from the statistical and risk-neutral Lévy measures. Specifically, we have that

$$\left(\frac{dQ}{dP}\right)_t = \exp\left\{-t\int_{-\infty}^{\infty}[Y(x)-1]k_P(x)\,dx\right\}\prod_{S\le t} Y[\Delta X(s)], \tag{22}$$

where $Y(x)$ is given by the equation

$$k_Q(x) = Y(x)k_P(x). \tag{23}$$

Hence, unlike the situation with diffusions, where options are redundant assets, option prices in a jump model can be used to infer the nature of the measure change process, provided, as noted earlier, that one restricts attention to approximating FA processes that exclude moves in a small interval about zero, say $(-\epsilon, \epsilon)$. Consequently, one can infer the prices of jump risks conditional on the size and sign of the jump. For the special case when the CGMY model describes the statistical and risk-neutral processes, we have that

$$Y(x) = \begin{cases} \dfrac{\tilde{C}}{C} x^{Y-\tilde{Y}} \exp\left(-(\tilde{M} - M)x\right) & x > \epsilon, \\ \\ \dfrac{\tilde{C}}{C} |x|^{Y-\tilde{Y}} \exp\left[-(\tilde{G} - G)|x|\right] & x < -\epsilon. \end{cases} \tag{24}$$

We shall comment further on the explicit form of this measure change in the light of our parameter estimates.

V. Data and Estimation Methodology

In an ideal context, one would obtain data on the time series of stock prices and the prices of options on the stock over a common time interval and then jointly evaluate the likelihood of observing this data on the assumption that the statistical and risk-neutral processes are parameterized by the extended CGYM class with parameters C, G, M, Y, η, $\tilde{C}$, $\tilde{G}$, $\tilde{M}$, $\tilde{Y}$, and $\tilde{\eta}$, along with the mean return μ of the statistical process. We estimate the statistical parameters from time-series data on the asset prices over the period January 1, 1994 to December 31, 1998. For the risk-neutral process, we follow the traditional practice established in the literature (See Bakshi, Cao, and Chen 1997) and estimate risk-neutral parameters on a set of days from closing option prices. We discuss the details of each of these two estimations separately in the following two subsections.

The data for both estimations were made available by Morgan Stanley Dean Witter and are made up of time series on 13 stock prices, with ticker symbols AMZN, BA, GE, HWP, IBM, INTC, JNJ, MCD, MMM, MRK, MSFT, WMT, and XON, and eight market indexes, with tickers BIX, BKX, DRG, RUT, SPX, SOX, XAU, and XOI. For the risk-neutral estimates, we employed closing option prices on five underlying assets, AMZN, IBM, INTC, MSFT, and SPX, for 5 midmonth Wednesdays (October 14, 1998; November 11, 1998; December 9, 1998; January 13, 1999; and February 10, 1999), with maturities between 1 and 2 months. The option prices are midmarket quotes for European options obtained first by determining volatility using a finite-difference, American option-pricing model calibrated to market American option prices where appropriate and then determining a European option price from this volatility estimate.

A. The Statistical Estimation and Results

For each underlying asset, we formed the time series of daily log price relatives and then estimated the parameters of the Lévy density C, G, M, Y, and η from the mean-adjusted return data. Direct maximum likelihood estimation is computationally expensive, as it requires a Fourier inversion for each data point to evaluate the density, and these inversions must be nested into a gradient search optimization algorithm for the parameter estimation.

The fast Fourier transform was used to invert the characteristic function once for each parameter setting. This method efficiently renders the level of the probability density at a prespecified set of values for returns. For integration spacing of .25, the density is obtained at a return spacing of $8\pi/N$, where N is a power of 2 used in the fast Fourier discrete transform. For $N = 4{,}096$, the return spacing is too coarse at .00613592. We used, instead, $N = 16{,}384$ and a return spacing of .00153398.

With the density evaluated at these prespecified points, we binned the return series by counting the number of observations at each prespecified return point, assigning data observations to the closest prespecified return point. We then searched for parameter estimates that maximized the likelihood of this binned data. The reported estimates are thus for this binned maximum likelihood estimation using the fast Fourier transform.

For the standard errors, we employ the inverse of the information matrix when the parameter estimates are in the interior of the parameter space. In the cases where the diffusion coefficient is estimated at the boundary of the parameter set at the value of zero, we provide the conditional standard errors of the other parameter estimates on inversion of the partial information matrix with respect to the other interior parameter estimates. To test the null hypothesis that the diffusion coefficient is zero, which is a test on the boundary of the parameter space, we employ a locally mean most powerful (LMMP) test statistic developed by King and Wu (1997). The statistic is normal with mean zero and unit variance under the null hypothesis and is reported when it is positive. It is based on the score function computed at the null.

The results of the estimates for 13 names and eight indices are presented in table 2 using both parameterizations, the implied VG parameters and the proper CGMY parameters. The estimation was conducted in the parameterization σ, ν, θ, η, and Y, with C, G, and M computed internally in accordance with equations (8), (9), and (10). In a few cases, standard errors were not available because of a lack of positive definiteness of the estimated information matrix.

The estimated densities have a variety of shapes, ranging from a diffusion component in MSFT to pure jump processes of infinite variation in the case of the index DRG. All of the indices are consistent with processes of infinite activity and finite variation. To appreciate further the range of possibilities, we present graphs of five of the fitted densities along with the empirical scatter of the binned data on daily log returns. First, we present the characteristic long necks of the SPX and RUT in figures 2 and 3.

We next present the bell-shaped structure in MSFT and XAU in figures 4 and 5.

Finally, we present a possible jump-diffusion case, as reflected in BA in figure 6.

B. The Risk-Neutral Estimation and Results

For each of the five underlying assets and for each of the 5 days, we obtained parameter estimates of the risk-neutral process by nonlinear least squares minimization of pricing errors from out-of-the-money closing option prices. For the computation of the model's option price, we followed Carr and Madan (1998) and inverted the analytical Fourier transform in log strike of the call prices dampened by an exponential factor. The results for the risk-neutral estimation are presented in table 3.

VI. Discussion of Results

We discuss the results from four perspectives. First, we consider the issue of skewness and kurtosis in returns. Next, we consider partitioning the total quadratic variation into its pure jump and diffusion components. We then address questions related to the fine structure of the process as embedded in the parameter Y. Finally, we close with a discussion of the nature of the implicit measure change.

A. Skewness and Kurtosis

The evidence on statistical skewness is mixed. Of the 20 estimations, θ is significantly negative in five cases, which include SPX and RUT. Computing the exact skewness using the moment equation (19), we find negative skewness under the historical measure for just IBM, RUT, and SPX, respectively, at levels $-.0461$, $-.0047$, and $-.0028$. In the remainder of the cases, skewness is zero for seven cases and slightly positive for the remaining 10 cases. The kurtosis is generally above 3, and the excess kurtosis is as large as .1758 for WMT, while it is substantial for INTC, where it is estimated at 16.19 when volatility is low at .02. The historical levels of volatility, skewness, and kurtosis, as computed by the moment equations, are reported in table 4.

In contrast, the risk-neutral process is definitely negatively skewed with M dominating G and θ negative in every case. The skewness as computed using the moment equations is negative in every case except for AMZN on January 13, when it is slightly positive. We also note that there is considerable variability in the skewness on our individual stocks across time, showing a general decline between October 1998 and February 2000. On the SPX, however, skewness is more stable across time.

The risk-neutral kurtosis is substantially higher than the historical levels

TABLE 2 **Results of Maximum Likelihood Estimation of the Binned Data on Continuously Compounded Daily Returns at a Return Spacing of .001534**

	σ	ν	θ	η	Y	C	G	M	LL/Z
BA	.2428	.0152	.0118	.0914	−.0719	65.65	47.38	46.98	3354.93
	(.1425)	(.0251)	(.0011)	(.1927)	(.3461)				.7036
GE	.1331	.0468	.0123	.0164	.0037	21.34	49.78	48.40	3551.80
	(.1348)	(.1359)	(.0007)	(.0929)	(.5842)				.1495
HWP	.2239	.0389	.0160	.0981	.0931	25.72	32.36	31.72	3048.21
	(.1227)	(.0652)	(.0014)	(.1730)	(.3621)				1.2088
IBM	.0706	.6655	−.1243	.0201	.7836	1.50	22.18	27.12	3262.24
	(1.0209)	(.9959)	(.0017)	(.0183)	(.2373)				1.7919
INTC	.6879	.0020	.0194	1e − 5	−.7904	4.94	45.74	45.66	2996.86
	(.6679)	(.0047)	(.0011)		(.6356)				NA
JNJ	.0424	.0951	0	.0113	.7515	10.52	108.06	108.06	3527.91
	(.0107)	(.0614)	(.0033)	(.0131)	(.1123)				.5925
MCD	.0162	12.14	5.2e − 8	7.9e − 4	1.50683	.08	25.04	25.04	3515.78
	(.0034)	(9.8020)	(.1097)	(.001)	(.1449)				NA
MMM	.1888	.0075	.0081	9e − 6	1.0023	133.77	86.86	86.41	3595.23
	(.0584)	(.0066)	(.0019)		(.1592)				NA
MRK	.1168	.0689	.0135	1.8e − 6	.1172	14.50	47.11	45.13	3434.08
	(.0982)	(.1758)	9e − 4	(.0623)	(.4745)				NA
MSFT	.2305	.0036	.0085	.3815	.1191	280.11	102.84	102.53	3112.5
	(.0479)	(.0027)	NA	(.3288)	NA				1.8599

WMT	.1905	.0422	.0161	.0268	−.0963	23.70	36.59	35.70	3236.39
	(.1891)	(.1149)	(7.8e − 4)	(.1463)	(.6482)				.1749
XON	.0709	.0255	−5.5e − 6	8.7e − 8	.4789	39.27	124.99	124.99	3664.43
	(.0199)	(.0202)	(.0166)		(.1314)				NA
BIX	.0189	3.1121	.0018	4.34e − 8	1.2341	.32	47.76	37.42	3710.07
	(.0023)	(.5491)	NA	.00003	NA				.1014
BKX	.1476	.0195	.0088	1.7e − 6	.0734	51.34	69.04	68.25	3680.18
	(.0743)	(.0288)	(.0010)		(.2599)				NA
DRG	.0255	.5729	−5.02e − 7	4.7e − 4	.9315	1.75	73.39	73.39	3872.48
	(.0048)	(.3387)	(.0044)		(.0995)				NA
RUT	.0597	.0678	−.0035	5.8e − 9	.3196	14.75	89.99	91.99	4401.39
	(.0218)	(.0672)	(5.4e − 4)		(.1831)				NA
SOX	.0271	2.2557	.00001	5.34e − 8	1.3814	.44	34.76	34.73	2731.49
	(.0048)	(.8011)	(.0129)	NA	(.0591)				NA
SPX	.0739	.0403	−.0042	1.9e − 10	.2495	24.79	94.45	95.79	4258.5
	(.0311)	(.0470)	(6.6e − 4)		(.2082)				NA
XAU	.1909	.0080	.0074	1.66e − 7	.3071	125.05	83.04	82.64	2732.58
	(.0266)	(.0046)	NA	NA	NA				NA
XOI	.1192	.0073	.0046	.0059	.0684	137.05	139.26	138.61	3112.54
	(.0473)	(.0071)	(.0013)	(.0814)	(.1917)				.0113

NOTE.—We report the VG parameter estimates σ, ν, and θ, and the transform to C, G, and M, as per equations (8), (9), and (10), along with the diffusion parameter η and the fine structure parameter Y. The final column reports the log likelihood and the LMMP Z-statistic where appropriate. Standard errors are in parentheses.

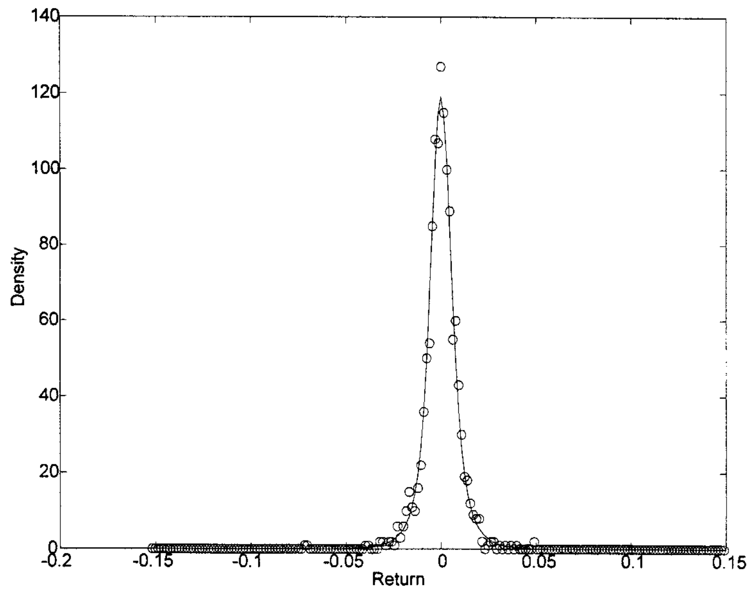

FIG. 2.—SPX MLE density fit

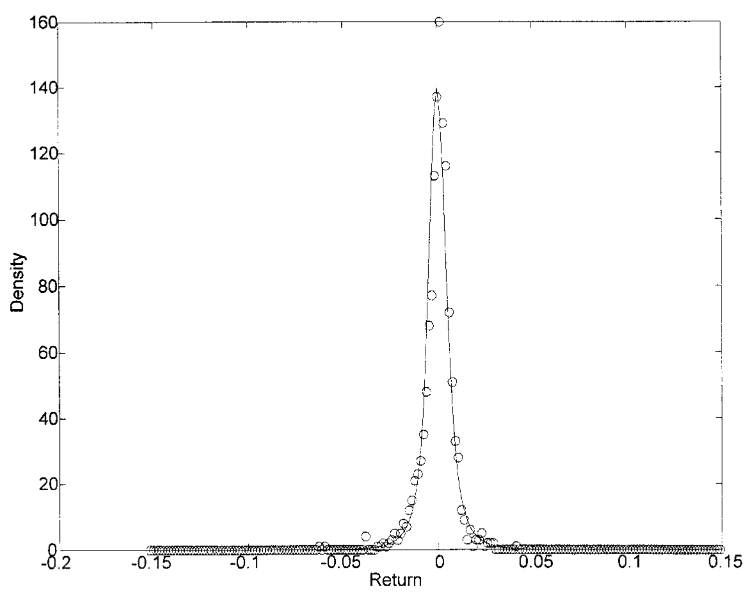

FIG. 3.—RUT MLE density fit

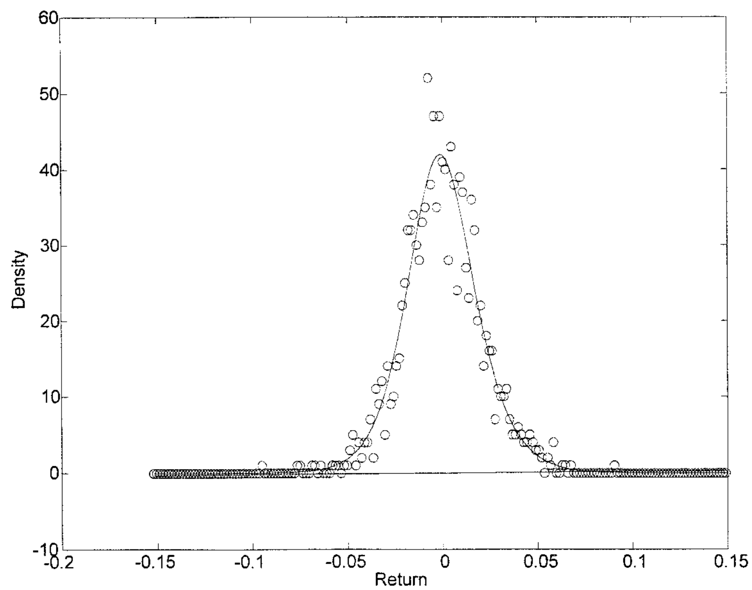

FIG. 4.—MSFT MLE density fit

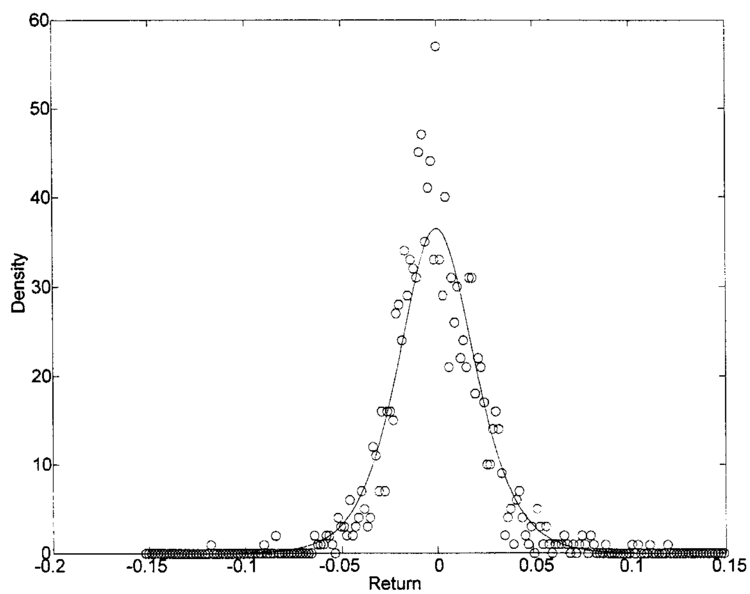

FIG. 5.—XAU MLE density fit

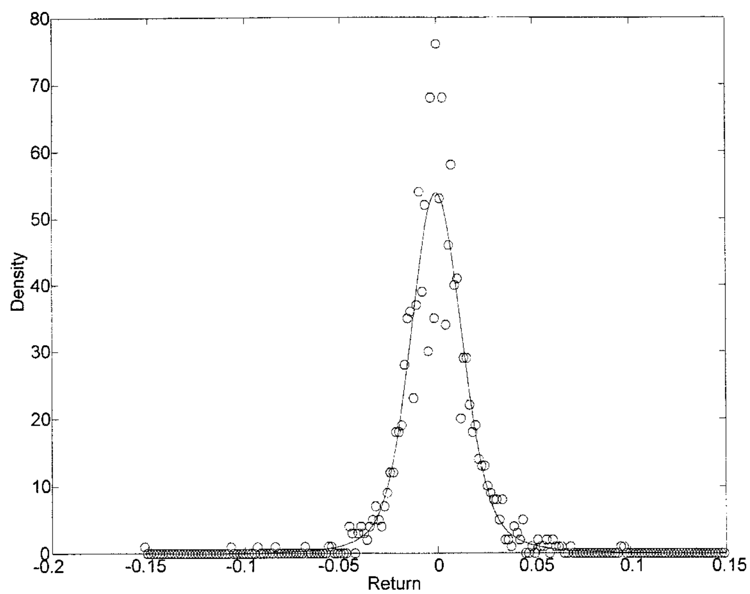

FIG. 6.—BA MLE density fit

for this statistic. On the SPX, excess kurtosis rises to 1.693 on November 11, while the historical level is just .0339. The risk-neutral higher moments are reported in table 5.

B. Decomposition of Quadratic Variation

A surprising feature of the results on the decomposition of quadratic variation is that, for all of the indices, the diffusion component is absent. On the individual stocks, the diffusion component is also absent for five companies and is positive but insignificant in the remaining seven cases. These are BA, GE, HWP, IBM, JNJ, MSFT, and WMT. We may employ (21) to determine the proportion of the total quadratic variation contributed by the diffusion component, and this is 15.32%, 1.48%, 12.60%, 0.71%, 0.23%, 62.29%, and 2.61% of the aggregate quadratic variation for BA, GE, HWP, IBM, JNJ, MSFT, and WMT, respectively.

A collective view of these results suggests that the diffusion components are diversifiable, while the systematic components, as reflected in the indices, are pure jump processes. This view is consistent with a single-index model in which the return distribution of the single factor is highly peaked near zero, to reflect long periods of little or no movement, coupled with fat tails, to reflect occasional movement of the whole market in one direction or the other. These findings also suggest that the diffusion components should be small in the risk-neutral process, as they can be costlessly diversified away.

To evaluate this conjecture, we took the parameter estimates for each stock

TABLE 3 **Results of Nonlinear Least Squares Estimation of Risk-Neutral Parameter Values on Five Underlying Assets for 5 Days for a Maturity of .1014**

	σ	ν	θ	η	Y	C	G	M	APE
spx1014	.3052	.1004	−.9558	.0312	−.0901	9.96	7.61	28.12	.3123
spx1111	.1616	.3277	−.3043	.0292	.1432	3.05	7.57	30.88	.5459
spx1209	.1775	.1704	−.4066	.0254	−.0008	5.86	10.31	36.09	.1324
spx01133	.1420	.2377	−.3657	.0326	.2227	4.21	9.18	45.43	.5379
spx0210	.1292	.0936	−.6990	.0189	.2155	10.69	13.21	97.00	.0632
amzn1014	.9015	.2193	−1.8315	.0684	.3072	4.60	1.78	6.29	.2687
amzn1111	.1207	.0447	−4.607	2e − 5	−.0069	22.39	4.82	637.80	.4514
amzn1209	.3999	.0589	−1.8308	.0712	.6442	16.97	7.08	29.97	.3431
amzn0113	1.3423	.0572	.1727	.0021	.2013	17.48	4.50	4.31	.6378
amzn0210	1.185	.0588	−2.141	.0016	−.001	17.02	3.63	6.68	.1498
ibm1014	.4105	.1430	−.6852	.0007	.0873	6.99	5.91	14.04	.0953
ibm1111	.0312	2.386	−.0938	.0428	1.0102	.42	4.37	191.20	.0358
ibm1209	.1157	.2740	−.2702	.0496	.4464	3.65	10.68	51.06	.0425
ibm0113	.3416	.1041	−.3816	.0458	.1430	9.61	9.97	16.51	.0628
ibm0210	.4344	.1083	−.3726	.0051	.0043	9.23	8.11	12.06	.3965
intc1014	.4072	.1022	−.7057	.0179	.0719	9.78	7.41	15.93	.0238
intc1111	.3517	.0277	.7767	.0003	.0004	36.06	18.67	31.23	.0268
intc1209	.4161	.0059	−1.739	.0026	.0006	170.04	35.40	55.48	.0610
intc0113	.1452	.1536	−.1497	.0869	.5757	6.51	18.75	32.95	.0379
intc0210	.4697	.0144	−1.132	.0437	.020	69.51	20.49	30.75	.0346
msft1014	.4669	.0495	−1.2142	.0631	.0483	20.20	9.14	20.28	.0813
msft1111	.4089	.0279	−.8034	.0041	−.1341	35.78	16.43	26.04	.0609
msft1209	.2242	.9087	−.0881	.0774	.4456	1.10	5.09	8.60	.0606
msft0113	.4757	.0246	−.7350	.0058	.0011	40.62	15.98	22.48	.0823
msft0210	.4383	.0269	−.7797	7e − 5	.0003	37.16	16.03	24.14	.0630

NOTE.—Only out-of-the-money options were used in the estimation.We report the VG parameters σ, ν, and θ; the transform to C, G, M as per eqq. (8), (9), and (10); and the diffusion parameter η and fine structure parameter Y. The final column reports the average pricing error in each case.

on 3 of the 5 days with the best fit in terms of average pricing error and computed the proportion of the total quadratic variation attributable to the diffusion component. We found that, in each case, the proportion of the quadratic variation of the risk-neutral process due to the diffusion component was zero. Hence, we tentatively conclude that diffusion components are not priced in the market for risks.

C. The Fine Structure of Returns

Regarding the fine structure of statistical returns, we find that, for just three of the individual stocks (BA, INTC, and WMT), the statistical jump component is one of finite activity. However, the null hypothesis of a VG process cannot be rejected for any of these cases. In all of the other cases, we have infinite activity, and, except for BIX, SOX, and MCD, we typically estimate a finite variation process. Thus, the jump component mainly reflects both infinite activity and finite variation for the statistical process.

With respect to the risk-neutral process, we note that essentially all of the processes are infinite-activity, finite-variation processes. This is reasonable in our view, as infinite variation comes from a high degree of activity near zero

TABLE 4 Statistical Levels of Volatility, Skewness, and Kurtosis as Computed Using the Moment Equations (18), (19), and (20)

	Volatility	Skewness	Kurtosis
BA	.2335	.0021	3.0444
GE	.1350	.0125	3.1344
HWP	.2763	.0055	3.0618
IBM	.2385	−.0461	3.0818
INTC	.0196	.0103	16.1960
JNJ	.2351	0	3.0043
MCD	.2458	0	3.0194
MMM	.1786	−.0001	3
MRK	.1430	.0177	3.1252
MSFT	.4834	.00006	3.0008
WMT	.1660	.0130	3.1758
XON	.2122	0	3.0055
BIX	.2102	.0186	3.0180
BKX	.1699	.0028	3.0415
DRG	.1849	0	3.0120
RUT	.1168	−.0047	3.0399
SOX	.3781	.00003	3.0058
SPX	.1253	−.0028	3.0339
XAU	.3583	.00038	3.0052
XOI	.1393	.0007	3.0151

and the pricing process is essentially pricing large moves with little attention to the small moves. These considerations are suggestive of finite variation in the risk-neutral process. We also observe that, in all the cases, both statistical and risk-neutral, the Lévy density is consistent with the hypothesis of complete monotonicity.

D. Explicit Measure Changes

For each of the four assets for which we have estimated both the risk-neutral and statistical CGMY jump components, we use equation (24) to explicitly construct the measure change function $Y(x)$ on an approximating finite activity process truncating small moves. We use, for each asset, the risk-neutral parameter values for 1 of the 5 days on which the parameters were estimated. Figure 7 presents the graph of the measure change function on the SPX for January 13, 1999.

We observe that the function rises on both sides, with a much steeper ascent on the left. This is indicative of risk premia for large jump sizes on both sides of zero. The picture is quite typical and is fairly consistently observed in the SPX market. A more symmetric U-shaped measure change is observed for MSFT on December 9, 1998, and this is shown in figure 8.

A somewhat different shape is observed for INTC, as shown in figure 9. This reflects significant premia for down moves but milder premium levels for up moves.

It is interesting to enquire into the reasons for the shape of the measure change function $Y(x)$. In a two-person equilibrium with heterogeneous beliefs and preferences, investors take a nonzero position in options, as shown, for

TABLE 5 **Risk Neutral Levels of Volatility, Skewness, and Kurtosis as Computed Using the Moment Equations**

	Volatility	Skewness	Kurtosis
spx1014	.3999	−.6297	3.6540
spx1111	.2706	−.8196	4.1693
spx1209	.2453	−.7068	3.8643
spx0113	.2849	−.6263	3.6746
spx0210	.3198	−.4077	3.2699
amzn1014	1.3218	−.6204	3.7415
amzn1111	.9779	−.4258	3.2716
amzn1209	1.1042	−.1463	3.0460
amzn0113	1.5042	.0165	3.1150
amzn0210	1.2928	−.2764	3.2289
ibm1014	.5186	−.4817	3.5148
ibm1111	.3197	−.6801	3.9704
ibm1209	.3026	−.4221	3.3411
ibm0113	.4219	−.2411	3.2452
ibm0210	.4532	−.2598	3.3669
intc1014	.4958	−.3812	3.3567
intc1111	.3749	−.1655	3.1017
intc1209	.4374	−.0679	3.0207
intc0113	.3688	−.1001	3.0539
intc0210	.5044	−.0922	3.0454
msft1014	.5688	−.2770	3.1813
msft1111	.3649	−.1978	3.1555
msft1209	.3442	−.4311	3.9836
msft0113	.4900	−.1086	3.0815
msft0210	.4567	−.1341	3.0928

example, in Franke, Stapleton, and Subrahmanyam (1998) or Carr and Madan (2001). Hence, one may infer the measure change if one has data on preferences and investor positions. It is well known that the measure change is given by the marginal utility of the position times the ratio of subjective to objective probabilities. Specifically, one may write that as

$$Y(x) = \frac{U'[c(Se^x)]p_S(x)}{U'[c(S)]p_O(x)},$$

where U is the investor utility function, $p_S(x)$ is the investor subjective probability of a jump of size x in the log of the stock price, $p_O(x)$ is the corresponding true statistical probability, and $c(Se^x)$ is the state contingent claim being held by the investor. For a Lucas representative agent holding the stock and under rational expectations, that is, $p_S(x) = p_O(x)$, we deduce that

$$Y(x) = \frac{U'(Se^x)}{U'(S)},$$

which is a function that is monotonically decreasing in x for all concave utility functions.

When markets are incomplete and beliefs are heterogeneous, one needs to combine preferences, positions, and beliefs more carefully in order to infer the nature of the function $Y(x)$. If we take the view that option writers have

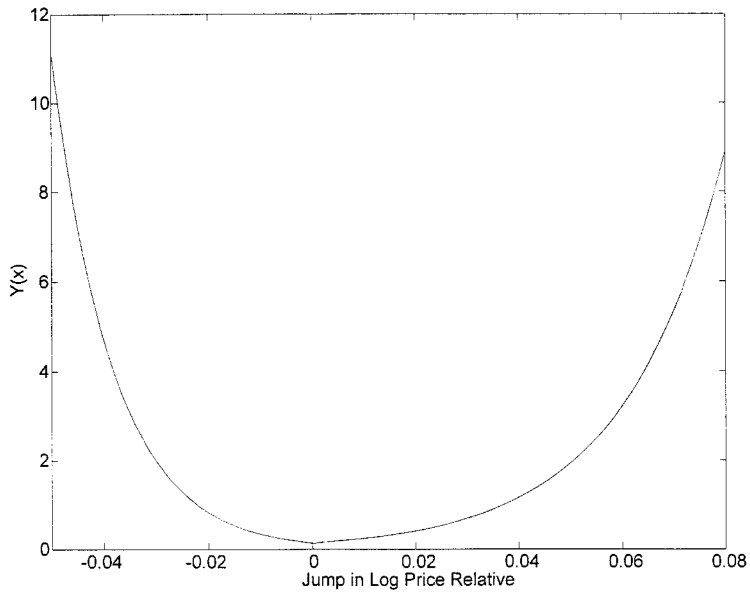

FIG. 7.—Measure change density for SPX, January 13, 1999

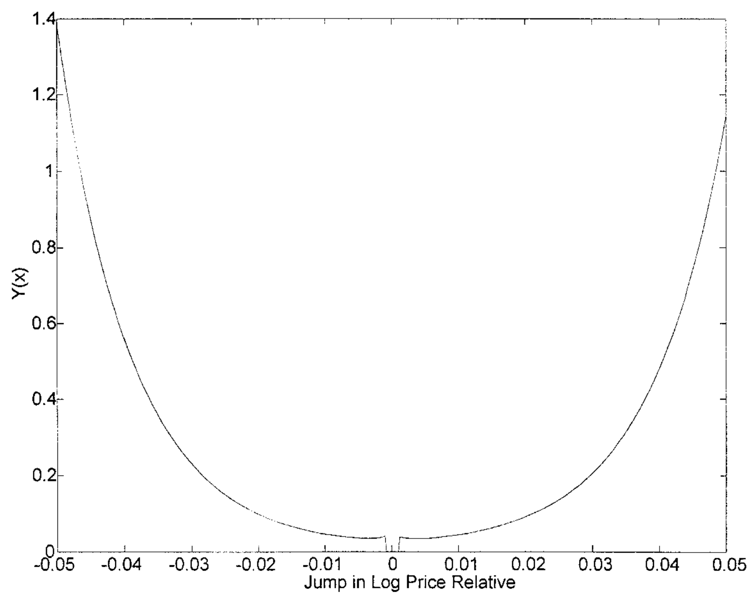

FIG. 8.—Measure change density for MSFT, December 9, 1998

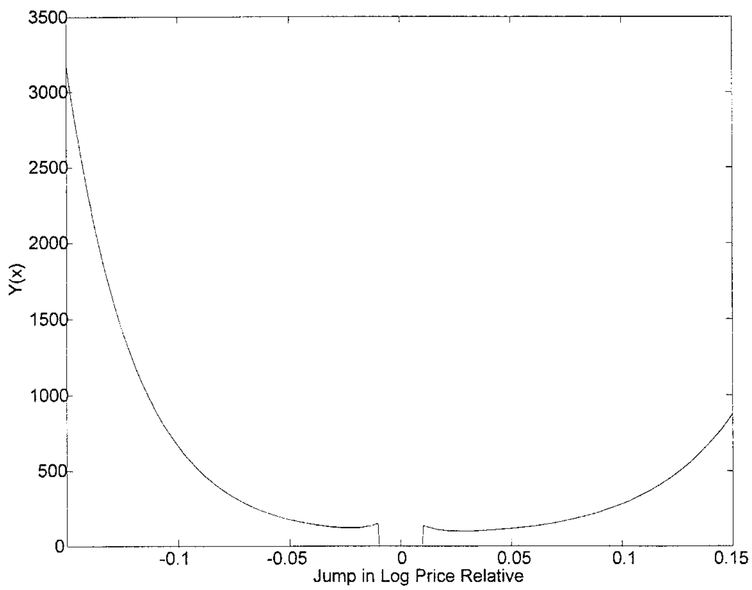

FIG. 9.—Measure change density for INTC, October 14, 1998

probability beliefs closest to the objective statistical probability, then, for individuals satisfying $p_S(x) = p_O(x)$, the position $c(Se^x) = g(x)$ is that of a delta-hedged option writer, with $g'(0) = 0$ and $g''(0) < 0$. The shape of $g(x)$ is that of an inverted U. It follows that

$$Y(x) = \frac{U'[g(x)]}{U'[g(0)]}$$

is of the form observed in our estimations. Furthermore, the relative rate of decrease of g on the two sides of zero is likely to be influenced by the structure of open interest in the market in put and call options. Hence, we conjecture that the structure of open interest in the market will be an important determinant of the shape of market risk premia as reflected in the measure change function $Y(x)$.

VII. Conclusions

This article generalizes the VG model to allow for Lévy processes with both a diffusion component and a Lévy measure that parametrically allows for processes with a finite or infinite activity, and with finite or infinite variation. The final model is termed the extended CGMY model, and we derive its characteristic function in closed form, which allows us to describe many of its properties.

The model is estimated on both time-series and option data, and it is observed that market indices lack a diffusion component. This leads to the conjecture that diffusion components observed in individual stock time series are diversified away in the index and, hence, the risk-neutral process should be devoid of a diffusion component. Estimation on option price data tends to provide confirmation of this conjecture.

We also report significantly greater skewness and kurtosis in the risk-neutral process than in the statistical process. We find that risk-neutral processes are mainly infinite-activity, finite-variation processes, while infinite variation may be prevalent in the statistical process for indices and for some stocks.

Broadly, our results suggest that option-pricing models should be built using completely monotone Lévy densities that integrate to infinity and are consistent with finite variation. We explicitly construct the embedded process for the measure change using approximating finite activity processes that exclude a small neighborhood of zero. Our results lead us to conjecture that the measure change process is related to the structure of open positions in the market.

References

Bakshi, G.; Cao, C.; and Chen, Z. 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52:2003–49.

Black, F., and Scholes, M. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81:637–54.

Carr, P., and Madan, D. B. 1998. Option valuation using the fast Fourier transform. *Journal of Computational Finance* 2:61–73.

Carr, P., and Madan, D. B. 2001. Optimal positioning in derivative securities. *Quantitative Finance* 1:19–37.

Cox J. C., and Ross, S. A. 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3:145–66.

Eberlein, E.; Keller, U.; and Prause, K. 1998. New insights into smile, mispricing, and value at risk: The hyperbolic model. *Journal of Business* 71:371–406.

Franke, G.; Stapleton, R.; and Subrahmanyam, M. 1998. Who buys and who sells options? The role of options in an economy with background risk. *Journal of Economic Theory* 82:89–109.

Geman, H.; Madan, D.; and Yor, M. 2001. Time changes for Lévy processes. *Mathematical Finance* 11:79–96..

King, M. L., and Wu, P. X. 1997. Locally optimal one-sided tests for multiparameter hypotheses. *Econometric Reviews* 16:131–56.

Koponen, I. 1995. Analytic approach to the problem of convergence of truncated Lévy flights towards the Gaussian stochastic process. *Physical Review E* 52:1197–99.

Jacod, J., and Shiryaev, A. 1980. *Limit Theorems for Stochastic Processes.* Berlin: Springer-Verlag.

Madan, D. B.; Carr, P.; and Chang, E. 1998. The variance gamma process and option pricing. *European Finance Review* 2:79–105.

Madan, D. B., and Milne, F. 1991. Option pricing with VG martingale components. *Mathematical Finance* 1:39–45.

Madan, D. B., and Seneta, E. 1990. The variance gamma (VG) model for share market returns. *Journal of Business* 63:511–24.

Merton, R. C. 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4:141–83.

Merton, R. C. 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3:125–44.

Vershik, A., and Yor, M. 1995. Multiplicativité du processus gamma et étude asymptotique des lois stables d'indice α, lorsque α tend vers 0. Prépublication. Paris: Laboratoire de Probabilités, Université Paris 6.