Modeling Financial Security Returns Using Lévy Processes^{*}

Liuren Wu[†]

Zicklin School of Business, Baruch College, City University of New York

(First draft March 17, 2004; this version December 5, 2005)

^{*}This is an early draft, with possibly many mistakes and incoherent arguments. I welcome comments, including references that I have inadvertently missed. I thank Vadim Linetsky (the editor), Peter Carr, Xiong Chen, Roger Lee, Haitao Li, and Hengyong Mo for comments. All remaining errors are mine.

[†]One Bernard Baruch Way, Box B10-225, New York, NY 10010-5585, USA; Tel.: +1-646-312-3509; fax: +1-646-312-3451; Email: liuren_wu@baruch.cuny.edu; homepage: <http://faculty.baruch.cuny.edu/lwu/>.

Modeling Financial Security Returns Using Lévy Processes

Abstract

Lévy processes can capture the behaviors of return innovations on a full range of financial securities. Applying stochastic time changes to the Lévy processes randomizes the clock on which the processes run, thus generating stochastic volatilities and stochastic higher return moments. Therefore, with appropriate choices of Lévy processes and stochastic time changes, we can capture the return dynamics of virtually all financial securities. Furthermore, in contrast to the hidden factor approach, we can readily assign explicit economic meanings to each Lévy process component and its associated time change in the return dynamics. The explicit economic mapping not only facilitates the interpretation of existing models and their structural parameters, but also adds economic intuition and direction for designing new models capturing new economic behaviors. Finally, under this framework, the analytical tractability of a model for derivative pricing and model estimation originates from the tractability of the Lévy process specification and the tractability of the activity rate dynamics underlying the time change. Thus, we can design tractable models using any combination of tractable Lévy specifications and tractable activity rate dynamics. In this regard, we can incorporate and therefore encompass all tractable models in literature into our framework as building blocks. Examples include Brownian motions, compound Poisson jumps, and other tractable jump specifications like variance gamma, dampened power law, normal inverse Gaussian, and so on for Lévy processes, and affine, quadratic, and $3/2$ processes for activity rate dynamics. In this chapter, I elaborate through examples on the generality of the framework in capturing the return behavior of virtually all financial securities, the explicit economic mapping that facilitates the interpretation and creation of new models capturing specific economic behaviors, and the tractability embedded in the framework for derivative pricing and model estimation.

JEL Classifications: G10, G12, G13.

Keywords: Lévy processes; return innovations; stochastic time changes; stochastic volatility; characteristic functions; exponential martingales; measure change; option pricing; Fourier inversion.

Contents

1	Introduction	1
2	Modeling return innovation distribution using Lévy processes	4
2.1	Lévy characteristics	5
2.2	Lévy examples	7
2.3	Empirical evidence	10
3	Generating stochastic volatility by applying stochastic time changes	12
3.1	Time changes and activity rates	13
3.2	Generating stochastic volatility from different economic sources	14
3.3	Theory and evidence on activity rate dynamics	16
4	Modeling financial security returns with time-changed Lévy processes	19
4.1	Constructing risk-neutral return dynamics	20
4.1.1	Examples	22
4.2	Market price of risks and statistical dynamics	25
4.3	More flexible market price of risk specifications	29
5	Option pricing under time-changed Lévy processes	33
5.1	Deriving the Fourier transforms	33
5.2	Computing the Fourier inversions	38
5.2.1	The cumulative distribution analogy	39

5.2.2	The probability density analogy	42
5.2.3	Fractional FFT	44
6	Estimating Lévy processes with and without time changes	45
6.1	Estimating statistical dynamic using time-series returns	45
6.2	Estimating risk-neutral dynamic to fit option prices	46
6.3	Static and dynamic consistency in model estimation	48
6.4	Joint estimation of statistical and risk-neutral dynamics	49
7	Concluding remarks	50

1. Introduction

Since Black and Scholes (1973), Brownian motion has emerged as the benchmark process for describing asset returns in continuous time. Brownian motion generates normally distributed return innovations. Merton (1976) augments the Brownian motion with a compound Poisson process with normally distributed jump sizes in the asset return. As a result, the return innovation distribution becomes a mixture of normals with Poisson probability weightings. These two innovation specifications have dominated the continuous-time finance literature for several decades, drawing criticisms that the continuous-time framework is not as flexible as the discrete-time framework: One can assume any arbitrary distribution for the return innovation in discrete time, but only normals or mixtures of normals could be generated from continuous-time specifications.

The recent advent of Lévy processes completely exonerates continuous-time finance from such criticisms. For virtually all distribution specifications, one can specify a Lévy process that generates such a distribution at a fixed horizon. While the Brownian motion component in a Lévy process generates a normal distribution, any non-normal distribution can be generated via the appropriate specification of the Lévy density for a Lévy jump process, which determines the arrival rate of jumps of all possible sizes.

Financial security returns can be driven by several economic forces. The impact of each force can vary stochastically over time. Accordingly, we can model the return innovation using several Lévy processes as building blocks matching the distributional behavior of shocks from different economic forces. Furthermore, applying stochastic time change to each Lévy component randomizes the clock on which the process runs, thus capturing the stochastically varying impacts from different economic forces. Statistically, applying stochastic time changes on different Lévy components can generate both stochastic volatility and stochastic higher return moments, both of which are well-documented features for financial securities. Therefore, with appropriate choices of Lévy processes and stochastic time changes, we can capture the return dynamics of virtually all financial securities.

Generality is not the only virtue of Lévy processes. By modeling return dynamics using different combinations of Lévy components with time changes, we can readily assign explicit economic meanings to each Lévy component and its associated time change in the return dynamics. The explicit

economic mapping not only facilitates the interpretation of existing models and their structural parameters, but also adds economic intuition and direction for designing new models that are parsimonious and yet adequate in capturing the requisite economic behaviors. In contrast, a common approach in the literature is to model returns by a set of hidden statistical factors. Factor rotations make it inherently difficult to assign economic meanings to the statistical factors. The absence of economic mapping also makes the model design process opaque. For One often finds that a generic hidden-factor model cannot match the requisite target behaviors of the financial securities returns, and yet many parameters of the model are difficult to identify empirically. The issue of being both “too little” in performance and “too much” in model identification can only be solved by exhaustive econometric analysis. By estimating different restricted versions of the factor models with increasing numbers of factors, one can hopefully identify which set of parameters are redundant and how many additional factors are needed.

The generality of the framework does not hinders its analytical tractability for derivative pricing and model estimation, either. When modeling return dynamics using Lévy processes with time changes, tractability of the return dynamics originates from tractability of the Lévy component specification and tractability of the activity rate dynamics underlying the time change. Thus, we can design tractable models using any combinations of tractable Lévy processes and tractable activity rate dynamics. In this regard, we can incorporate and hence *encompass* all tractable models in the literature as building blocks. Examples of tractable Lévy specifications include Brownian motions, compound Poisson jumps, and other tractable jump specifications like variance gamma, dampened power law, normal inverse Gaussian, and so on. Examples of tractable activity rate dynamics include the affine class of Duffie, Pan, and Singleton (2000), the quadratic class of Leippold and Wu (2003), and the $3/2$ process of Heston (1997) and Lewis (2001). By modeling financial securities returns with time-changed Lévy processes, we encompass all these models into one general and yet tractable framework.

Through examples, I elaborate in this chapter the three key virtues of Lévy processes with stochastic time changes: (i) the generality of the framework in capturing the return behavior of virtually all financial securities, (ii) the explicit economic mapping that facilitates the interpretation and creation of new models capturing specific economic behaviors, and (iii) the tractability embedded in the framework for derivative pricing and model estimation.

In designing models for a financial security return, the literature often starts by specifying a very general process with a set of hidden factors and then testing different restrictions on this general process. In this chapter, I take the opposite approach. First, I look at the data and identify stylized features that a reasonable model needs to capture. Second, I design different components of the model to match different features of the data and capture the impacts from different economic forces. The final step is to assemble all the parts together. Using time-changed Lévy processes matches this procedure well. First, we can choose Lévy components to match the properties of return innovations generated from different economic forces. Statistically, we ask the following set of questions: Do we need a continuous component? Do we need a jump component? Do the jumps arrive frequently or are they rare but large events? Do up and down movements show different behaviors?

Once we have chosen the appropriate Lévy components, we can use time changes to capture the intensity variation for the different components and generate stochastic volatilities and stochastic higher return moments from different economic sources. We use time changes to address the following questions: Is stochastic volatility driven by intensity variations of small movements (Brownian motion) or large movements (jumps)? Do the intensities of different types of movements vary synchronously or separately? Do they show any dynamic interactions? Based on answers to these questions, we can apply different time changes to different Lévy components and model their intensity dynamics in a way matching their observed dynamic interactions.

The final step involves assembling the different Lévy components with or without time changes together into the asset return dynamics. When the dynamics are specified under the risk-neutral measure for derivative pricing, adjustments are necessary to guarantee the martingale property.

When designing models, tractability requirement often comes from derivative pricing when we need to take expectations of future contingent payoffs under the risk-neutral measure to obtain its present value. Thus, it is often convenient to start by specifying a tractable return dynamics under the risk-neutral measure. Then the statistical dynamics can be derived based on market price of risk specifications. The less stringent requirement for tractability for the statistical dynamics often allows us to specify very flexible market price of risk specifications, with the constraints only coming from reasonability for investor behaviors and parsimony for econometric identification.

In designing the different Lévy components and applying the time changes, I quote Albert Einstein as the guiding principle: “*Everything should be made as simple as possible, but not simpler.*” The explicit economic purpose for each Lévy component and its time change allows us to abide by this guiding principle much more easily than in a general hidden statistical factor framework.

The rest of the chapter is organized as follows. The next section discusses Lévy processes and how they can be used to model return innovations. Section 3 discusses how to use time changes to generate stochastic volatility and stochastic skewness from different sources. Section 4 discusses how to assemble different pieces together, how to satisfy the martingale condition under the risk-neutral measure, and how to derive the statistical dynamics based on market price of risk specifications. Section 5 discusses option pricing under time-changed Lévy processes. Section 6 addresses the estimation issues using time-series returns and/or option prices. Section 7 concludes.

2. Modeling return innovation distribution using Lévy processes

A Lévy process is a continuous time stochastic process with stationary independent increments, analogous to iid innovations in a discrete-time setting. Until very recently, the finance literature narrowly focuses on two examples of Lévy processes: the Brownian motion underlying the Black and Scholes (1973) model and the compound Poisson process with normal jump sizes underlying the jump diffusion model of Merton (1976). A Brownian motion generates normal innovations. The compound Poisson process in the Merton model generates return non-normality through a mixture of normal distributions with Poisson probability weightings. A general Lévy process can generate a much wider range of distributional behaviors through different types of jump specifications. The compound Poisson process used in the Merton model generates a finite number of jumps within a finite time interval. Such a jump process is suitable to capture rare and large events such as market crashes and corporate defaults. Nevertheless, many observe that asset prices can also display many small jumps on a fine time scale. A general Lévy process can not only generate continuous movements via a Brownian motion and rare and large events via a compound Poisson process, but it can also generate frequent jumps of different sizes.

2.1. Lévy characteristics

We start with a one-dimensional real-valued stochastic process $\{X_t | t \geq 0\}$ with $X_0 = 0$ defined on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ endowed with a standard complete filtration $\mathbf{F} = \{\mathcal{F}_t | t \geq 0\}$. We assume that X is a Lévy process with respect to the filtration $\mathbf{F}$, that is, X_t is adapted to $\mathcal{F}_t$, the sample paths of X are right-continuous with left limits, and $X_u - X_t$ is independent of $\mathcal{F}_t$ and distributed as X_{u-t} for $0 \leq t < u$. By the Lévy-Khintchine Theorem, the characteristic function of X_t has the form,

$$\phi_{X_t}(u) \equiv \mathbb{E}^{\mathbb{P}} [e^{iuX_t}] = e^{-t\psi_x(u)}, \quad t \geq 0, \quad (1)$$

where the *characteristic exponent* $\psi_x(u)$ is given by,

$$\psi_x(u) = -iu\mu + \frac{1}{2}u^2\sigma^2 + \int_{\mathbb{R}_0} (1 - e^{iux} + iux1_{|x|<1}) \pi(x)dx, \quad (2)$$

where $\mu \in \mathbb{R}$ describes the constant drift, $\sigma^2 \in \mathbb{R}^+$ describes the constant variance of the continuous component of the Lévy process, and the *Lévy density* $\pi(x)$ describes the arrival rates for jumps of every possible size x . The triplet (μ, σ^2, π) fully specifies the Lévy process X_t and is referred to as the *Lévy characteristics* (Bertoin (1996)).

With a fixed time horizon, any return distribution can be represented uniquely by its characteristic function and hence its characteristic exponent. Equation (2) illustrates that a Lévy process can generate a wide range of characteristic exponent behaviors through a flexible specification of the Lévy density $\pi(x)$.

The Lévy density $\pi(x)$ is defined on the real line excluding zero, $\mathbb{R}_0$. The truncation function $x1_{|x|<1}$ equals x when $|x| < 1$ and zero otherwise. Other truncation functions are also used in the literature as long as they are bounded, with compact support, and satisfy $h(x) = x$ in a neighborhood of zero (Jacod and Shiryaev (1987)).¹ The purpose of the truncation function is to analyze the jump properties around the singular point of zero jump size.

¹Commonly used truncation functions include $h(x) = x/(1+x^2)$, and $h(x) = 1 \wedge |x|$ (the minimum of 1 and $|x|$).

The characteristic function in (1) is defined on the real line $u \in \mathbb{R}$. In many applications, it is convenient to extend the definition to the complex plane, $u \in \mathcal{D} \subseteq \mathbb{C}$, where the characteristic exponent is well-defined. When $\phi_{X_t}(u)$ is defined on the complex plane, it is referred to as the *generalized Fourier transform* (Titchmarsh (1986)). It is also helpful to define the cumulant exponent of a Lévy process X_t ,

$$\phi_x(s) \equiv \frac{1}{t} \ln \mathbb{E} [e^{sX_t}] = s\mu + \frac{1}{2}s^2\sigma^2 + \int_{\mathbb{R}_0} (e^{sx} - 1 - sx1_{|x|<1}) \pi(x)dx, \quad s \in \mathcal{D}_s \subseteq \mathbb{C}, \quad (3)$$

where $\mathcal{D}_s$ denotes the subset of the complex plain under which the cumulant exponent is well-defined. Our extensions on the domains of the characteristic coefficient u and cumulant coefficient s implies that $\psi_x(u) = -\phi_x(iu)$ whenever the two are well-defined. Option pricing and likelihood estimation for Lévy processes often rely on the tractability of the characteristic exponent and specifically, analytical solutions to the integral in equation (2) or (3).

The sample paths of a pure jump Lévy process exhibit *finite activity* when the integral of the Lévy density is finite:

$$\int_{\mathbb{R}_0} \pi(x)dx = \lambda < \infty, \quad (4)$$

where λ measures the mean arrival rate of jumps. A finite activity jump process generates a finite number of jumps within any finite time interval.

When the integral in (4) is infinite, the sample paths exhibit *infinite activity*, and generate an infinite number of jumps within any finite interval. Nevertheless, the sample paths show *finite variation* if the following integral is finite:

$$\int_{\mathbb{R}_0} (|x| \wedge 1) \pi(x)dx < \infty. \quad (5)$$

When the integral in (5) is infinite, the jump process exhibit *infinite variation*, a property also shared by the Brownian motion. The truncation function in the definition of characteristic exponent is needed only for infinite variation jumps. When the integral in (5) is not finite, the sum of small jumps does not converge, but the sum of the jumps compensated by their mean converges. This special behavior generates the necessity for the truncation term in (2).

For all jump specifications, we require that the process exhibit *finite quadratic variation*:

$$\int_{\mathbb{R}_0} (1 \wedge x^2) \pi(x) dx < \infty, \quad (6)$$

a necessary condition for the jump process to be a semimartingale.

2.2. Lévy examples

Black and Scholes (1973) model the asset return by a purely continuous Lévy process and hence with $\pi(x) = 0$ for all x . The characteristic exponent is simply:

$$\psi(u) = -iu\mu + \frac{1}{2}u^2\sigma^2. \quad (7)$$

The associated normal probability density function is also well-known.

Merton (1976) incorporates an additional compound Poisson jump component with mean arrival rate λ . The jump size in the log asset return is normally distributed with mean μ_J and variance v_J , conditional on one jump occurring. The Lévy density of this jump component can be written as,

$$\pi(x) = \lambda \frac{1}{\sqrt{2\pi v_J}} \exp\left(-\frac{(x - \mu_J)^2}{2v_J}\right). \quad (8)$$

The characteristic exponent for this compound Poisson jump is:

$$\psi(u) = \lambda \left(1 - e^{iu\mu_J - \frac{1}{2}u^2 v_J}\right). \quad (9)$$

A key property of compound Poisson jumps is that the sample paths exhibit finite activity. Finite-activity jumps are useful in capturing large but rare events. For example, the credit-risk literature has used Poisson process extensively to model the random arrival of default events (Lando (1998), Duffie and Singleton (1999, 2003), and Duffie, Pedersen, and Singleton (2003)). More recently, Carr and Wu (2005b) use a Poisson jump with zero recovery to model the impact of corporate default on the stock price. Upon arrival, the stock price jumps to zero. Carr and Wu (2005a) use a Poisson jump with

random recovery to model the impact of sovereign default on its home currency price. Upon arrival, the currency price jumps down by a random amount.

Within the compound Poisson jump type, Kou (2002) proposes a double-exponential conditional distribution for the jump size. The Lévy density is given by,

$$\pi(x) = \begin{cases} \lambda\beta_+ \exp(-\beta_+x), & x > 0, \\ \lambda\beta_- \exp(-\beta_-|x|), & x < 0, \end{cases} \quad \lambda, \beta_+, \beta_- > 0. \quad (10)$$

Under this specification, the jump arrival rate increases monotonically with decreasing jump size. Asymmetry between up and down jumps are induced by the different exponential coefficients β_+ and β_- . The characteristic exponent for this pure jump process is,

$$\psi(u) = -\lambda \left[(\beta_+ - iu)^{-1} - \beta_+^{-1} + (\beta_- + iu)^{-1} - \beta_-^{-1} \right]. \quad (11)$$

Kou and Wang (2004) show that the double-exponential jump specification allows tractable pricing for American and some path-dependent options.

Although it is appropriate to use compound Poisson jumps to capture rare and large events such as market crashes and corporate defaults, many observe that asset prices actually display many small jumps. These types of behaviors are better captured by infinite-activity jumps, which generate infinite number of jumps within any finite time interval. A popular example that can generate different jump types is the CGMY model of Carr, Geman, Madan, and Yor (2002), with the following Lévy density,

$$\pi(x) = \begin{cases} \lambda\beta_+ \exp(-\beta_+x) x^{-\alpha-1}, & x > 0, \\ \lambda\beta_- \exp(-\beta_-|x|) |x|^{-\alpha-1}, & x < 0, \end{cases} \quad \lambda, \beta_+, \beta_- > 0, \alpha \in [-1, 2]. \quad (12)$$

In this specification, the power coefficient α controls the arrival frequency of small jumps and hence the jump type. With the power coefficient $\alpha = -1$, the Lévy density becomes the double-exponential specification in (10), the sample paths of which show finite activity. The model generates finite-activity jumps as long as $\alpha < 0$. When $\alpha \in [0, 1)$, the model generates jumps with infinite activity but finite variation. The jump process exhibits infinite variation when $\alpha \in [1, 2]$. The condition $\alpha \leq 2$ is necessary to guarantee finite quadratic variation. With $\alpha < 0$, the power term makes the jump arrival approaches

infinity as the jump size approaches zero. The larger the power coefficient, the higher the frequency of small jumps. The two exponential coefficients β_+ and β_- control the arrival of large jumps. The difference in the two coefficients generates asymmetry in the tails of the distribution.

The physics literature often refers to the specification in (12) as truncated Lévy flights. The CGMY terminology comes from the initials of the four authors in Carr, Geman, Madan, and Yor (2002), who regard the model as an extension of the variance gamma (VG) model of Madan and Seneta (1990) and Madan, Carr, and Chang (1998). Under the VG model, $\alpha = 0$. Wu (2006) label the specification in (12) as *exponentially dampened power law* (DPL), regarding it as the Lévy density of an α -stable Lévy process with exponential dampening. Wu shows that applying measure changes using exponential martingales to an α -stable Lévy process generates the exponentially dampened power law. Hence, the whole class of α -stable processes, made popular to the finance field by Mandelbrot (1963) and Fama (1965), can be regarded as a special class of the dampened power law.

When $\alpha \neq 0$ and $\alpha \neq 1$, the characteristic exponent associated with the dampened power law Lévy density specification takes the following form:

$$\psi(u) = -\Gamma(-\alpha)\lambda [(\beta_+ - iu)^\alpha - \beta_+^\alpha + (\beta_- + iu)^\alpha - \beta_-^\alpha] - iuC(h), \quad (13)$$

where $\Gamma(a) \equiv \int_0^\infty x^{a-1}e^{-x}dx$ is the gamma function and the linear term $C(h)$ is induced by the inclusion of a truncation function $h(x)$ for infinite-variation jumps when $\alpha > 1$. As I will make clear in later sections, in modeling return dynamics, any linear drift term in X_t will be canceled out by the corresponding term in its concavity adjustment. Hence, the exact form of the truncation function and the resultant coefficient $C(h)$ are immaterial for modeling and estimation. Wu (2006) explicitly carries out the integral in (3) through an expansion method and solves the truncation-induced term $C(h)$ under the truncation function $h(x) = x1_{|x|<1}$:

$$C(h) = \lambda(\beta_+(\Gamma(-\alpha)\alpha + \Gamma(1-\alpha, \beta_+)) - \beta_-(\Gamma(-\alpha)\alpha + \Gamma(1-\alpha, \beta_-))), \quad \alpha > 1, \quad (14)$$

where $\Gamma(a, b) \equiv \int_b^\infty x^{a-1}e^{-x}dx$ is the incomplete gamma function.

The dampened power law specification has two singular points at $\alpha = 0$ and $\alpha = 1$, under which the characteristic exponent takes different forms. The case of $\alpha = 0$ corresponds to the variance gamma model. Its characteristic exponent is,

$$\psi(u) = \lambda \ln(1 - iu/\beta_+) (1 + iu/\beta_-) = \lambda (\ln(\beta_+ - iu) - \ln \beta_+ + \ln(\beta_- + iu) - \ln \beta_-). \quad (15)$$

Since this process exhibits finite variation, we can perform the integral in (2) without the truncation function. When $\alpha = 1$, the characteristic exponent is (Wu (2006)),

$$\psi(u) = -\lambda ((\beta_+ - iu) \ln(\beta_+ - iu) / \beta_+ + \lambda (\beta_- + iu) \ln(\beta_- + iu) / \beta_-) - iu C(h), \quad (16)$$

where the truncation-induced term is given by $C(h) = \lambda (\Gamma(0, \beta_+) - \Gamma(0, \beta_-))$ under the truncation function $h(x) = x 1_{|x| < 1}$.

Other popular infinite-activity pure jump Lévy processes include the normal inverse Gaussian (NIG) process (Barndorff-Nielsen (1998)), the generalized hyperbolic process (Eberlein, Keller, and Prause (1998)), and the Meixner process (Schoutens (2003)). These processes all have tractable characteristic exponents.

2.3. Empirical evidence

Merton (1976)'s compound Poisson jump specification is suitable to capture large and rare events such as market crashes and corporate defaults. Nevertheless, recent empirical evidence suggests that infinite-activity jump specifications that generate frequent jumps of all sizes are better suited to capture the daily market movements of many financial securities such as stocks, stock indexes, and exchange rates. Furthermore, the distinction between continuous and discontinuous market movements is not at all clear cut. Instead, we observe movements of all sizes, with small movements arriving more frequently than large movements. This type of behavior asks for a Lévy density that is monotone in the absolute jump magnitude. The dampened power law specification in (12) has this monotonic behavior. Furthermore, when the power coefficient $\alpha \geq 1$, the arrival rate of small jumps is so frequent that the specification generates sample paths with infinite variation, a property also shared by the Brownian

motion. Hence, a Lévy process with infinite-variation jump provides a smooth transition from large jumps to small jumps and then to the continuous movements captured by a Brownian motion.

Several studies show that infinite-activity jumps perform better than finite-activity jumps in describing the statistical behavior of stock and stock index returns. Likelihood estimation of the dampened power law in Carr, Geman, Madan, and Yor (2002) on individual stocks and stock indexes shows that the estimates for the power coefficient α are mostly greater than zero. Li, Wells, and Yu (2004) use Markov Chain Monte Carlo method to estimate three Lévy specifications with stochastic time changes on the stock index. They find that infinite-activity jump specifications generates the better performance in capturing the index behaviors than finite-activity jumps.

Empirical studies using options show that using infinite-activity jumps also generate better option pricing performance. Carr and Wu (2003a) test the option pricing performance on the Merton jump-diffusion model, the variance gamma model, and their infinite-variation finite-moment log stable model. The pricing performance of the log stable model is the best among the three jump specifications. Huang and Wu (2004) apply various time changes on the three jump specifications to generate stochastic volatilities. They find that under all stochastic volatility specification, infinite-activity jumps perform significantly better than finite-activity jumps in pricing options.

Wu (2006) estimate the dampened power law using both the time-series returns and option prices on S&P 500 index. He obtains an estimate of the power coefficient at about 1.5. He also finds that although the exponential coefficient on down jumps β_- is large under the statistical measure, the estimate on its risk-neutral counterpart is not significantly different from zero. Without exponential dampening on down jumps, the return variance is infinite under the risk-neutral measure, even though it is finite under the statistical measure. As a result, the classic central limit theorem does not apply under the risk-neutral measure although it is applicable under the statistical measure. The difference under the two measures explains the empirical observation that the non-normality in the time-series index returns dissipates rapidly with time aggregation, but the risk-neutral return non-normality inferred from the options data persists to long option maturities.

When earlier studies use the compound Poisson jump to capture rate and large price movements, it is imperative to add a diffusion component to fill the gaps in between the arrival of the jumps.

However, if we start with an infinite-activity jump that can generate an infinite number of small and large movements within any finite interval, it is not clear that we still need a diffusion component to fill the gaps. Carr, Geman, Madan, and Yor (2002) conclude from their empirical study that a diffusion component is no longer necessary as long as they adopt an infinite activity pure jump process. Carr and Wu (2003a) arrive at similar conclusions in their infinite variation log stable model. Huang and Wu (2004) find that a diffusion return component is useful in their time-changed Lévy process setting in generating correlations with the diffusive activity rate process. Nevertheless, it is not clear whether the diffusion return component is still needed if the activity rate also follows a pure jump process and correlations are constructed through synchronous jumps in return and the activity rate.

Carr and Wu (2003b) identify the presence of jump and diffusion components in the underlying asset price process by investigating the short-maturity behavior of at-the-money and out-of-the-money options written on this asset. They prove that a jump component, if present, dominates the short-maturity behavior of out-of-the-money options and hence can readily be identified. A diffusion component, if present, usually dominates the short-maturity behavior of at-the-money options. However, an infinite-variation jump component can generate short-maturity behavior for at-the-money options that are similar to those generated from a diffusion process. The similar behavior make the identification of a diffusion component more difficult when an infinite-variation jump component is present. Nevertheless, Aït-Sahalia (2004) shows in a simple Lévy setting that when a diffusion component is present, the diffusion variance can be effectively identified even in the presence of infinite-variation jumps.

3. Generating stochastic volatility by applying stochastic time changes

It is well-documented that asset return volatilities are stochastic (Engle (2004)). Recent evidence from the derivatives market suggests that higher return moments such as skewness also vary significantly over time (David and Veronesi (1999), Johnson (2002), and Carr and Wu (2004a)). A convenient approach to generating stochastic volatility on non-normal return innovations is to apply stochastic time changes to a Lévy process; and a tractable way of generating stochastic skewness is to apply separate time changes to multiple Lévy components with different degrees of skewness. The random time change amounts to stochastically altering the clock on which the Lévy process is run. Intuitively, a

time change can be used to regulate the number of order arrivals that occur in a given time interval. More order arrivals generate higher return volatility (Ané and Geman (2000)). It can also be used to randomize the shocks from different economic sources. Separate time changes on different Lévy components can capture separate variations of different economic shocks.

3.1. Time changes and activity rates

Let X_t denote a Lévy process and let $t \rightarrow Z_t (t \geq 0)$ be an increasing right-continuous process with left limits that satisfy the usual technical conditions, we can define a new process Y obtained by evaluating X at Z , i.e.,

$$Y_t \equiv X_{Z_t}, \quad t \geq 0. \quad (17)$$

Monroe (1978) proves that every semimartingale can be written as a time-changed Brownian motion. Hence, equation in (17) is a very general specification. In principle, the random time Z_t can be modeled as a nondecreasing semimartingale,

$$Z_t = \mathcal{T}_t + \int_0^t \int_0^\infty x \mu(dt, dx), \quad (18)$$

where $\mathcal{T}_t$ is the locally deterministic and continuous component and $\mu(dt, dy)$ denotes the counting measure of the possible jumps of the semimartingale. The two components can be used to play different roles. Applying a time change defined by the positive jump component $\int_0^t \int_0^\infty x \mu(dt, dx)$ to a Brownian motion generates a new discontinuous process. If we model the positive jump component by a Lévy process, it is often referred to as a *Lévy subordinator*. A Lévy process subordinated by a Lévy subordinator yields a new Lévy process (Sato (1999)). Therefore, this component can be used to randomize the original return innovation defined by X to generate a refined return innovation distribution. For example, Madan and Seneta (1990) generate the variance-gamma pure jump Lévy process by applying a gamma time change to a Brownian motion.

To generate stochastic volatility on non-normal return innovations, I start directly with a Lévy process that *already* captures the non-normal return innovation distribution, and then apply a locally deterministic time change $\mathcal{T}_t$ purely for the purpose of generating stochastic volatilities and stochastic

higher return moments. We can characterize the locally deterministic time change in terms of its local intensity $v(t)$.

$$\mathcal{T}_t = \int_0^t v(u_-) du.$$

Carr and Wu (2004b) label $v(t)$ as the *instantaneous activity rate*. When X_t is a standard Brownian motion, v_t becomes the instantaneous variance of the Brownian motion. When X_t is a pure jump Lévy process, such as the compound Poisson jump process of Merton (1976), $v(t)$ is proportional to the jump arrival rate.

Although $\mathcal{T}_t$ is locally deterministic and continuous, the instantaneous activity rate process $v(t)$ can be fully stochastic and can jump. Given any continuous or discontinuous dynamics for $v(t)$, the integration over its sample path makes $\mathcal{T}_t$ locally predictable and continuous. Nevertheless, for $\mathcal{T}_t$ to be non-decreasing, the activity rate needs to be nonnegative, a natural requirement for diffusion variance and jump arrival rates.

3.2. *Generating stochastic volatility from different economic sources*

By applying stochastic time changes to Lévy processes, it becomes obvious that stochastic volatility can come from multiple sources. It can come from the instantaneous variance of a diffusion return component, or the arrival rate of a jump component, or both. Huang and Wu (2004) design and estimate a class of models for S&P 500 index returns based on the time-changed Lévy process framework. They allow the return innovation to contain both a diffusion component and a jump component. Then, they consider several cases where they apply stochastic time changes to (1) the diffusion component only (SV1), (2) the jump component only (SV2), (3) both components with one joint activity rate (SV3), and (4) both components with separate activity rates for each component (SV4). They find that by allowing the diffusion variance rate and the jump arrival rate to follow separate dynamic processes, the SV4 specification outperforms all the other single activity rate specifications in pricing the index options.

Applying separate stochastic time changes to different Lévy components also proves to be a tractable way of generating stochastic higher return moments such as skewness. In the SV4 specification of Huang and Wu (2004), one activity rate controls the intensity of a diffusion and hence a normal innovation component and the other activity rate controls the intensity of a negatively skewed pure jump

component. The variation of the two activity rates over time generates variation in the relative proportion of the diffusion versus the negatively-skewed jump return innovation components. As a result, the degree of the negative skewness for the index return varies over time (David and Veronesi (1999)).

Carr and Wu (2005b) apply the time-changed Lévy process framework to jointly price stock options and credit default swaps written on the same company. They assume that corporate default arrives via a Poisson process with stochastic arrival rate. Upon default, the stock price jumps to zero. Prior to default, the stock price follows a purely continuous process with stochastic volatility. Hence, the model decomposes the stock return into two Lévy components: (i) the continuous component that captures the market risk, and (ii) the jump component that captures the impact of credit risk. Separate time changes on the two components generate stochastic volatility for market movements and stochastic arrival for corporate default, respectively. Carr and Wu (2005a) use a similar specification to capture the correlation between sovereign credit default swap spreads and currency options. They assume that sovereign default induces a negative but random jump in the price of the home currency.

For stock indexes and the dollar (or euro) prices of emerging market currencies, the risk-neutral return distribution skewness is time-varying, but the sign stays negative across most of the sample period.² In contrast, for the exchange rates between two relatively symmetric economies, Carr and Wu (2004a) find that the risk-neutral currency return distribution inferred from option prices shows skewness that not only varies significantly over time in magnitudes, but also switches signs. To capture the stochastic skewness with possible sign switches, they decompose the currency return into two Lévy components that generate positive and negative skewness, respectively. Then, they apply separate stochastic time changes to the two Lévy components so that the relative proportion of the two components and hence the relative degree and direction of the return skewness can vary over time. They model the positively-skewed Lévy process with a jump component that only jumps upward and the negatively-skewed Lévy process with a jump component that only jumps downward. Furthermore, each process contains a diffusion component that is correlated with their respective activity rate process. The correlation is positive for the positive-skewed Lévy process and negative for the negative-skewed Lévy component. Thus, the up and down jumps generate short-term positive and negative skewness for

²See the evidence in David and Veronesi (1999) and Foresi and Wu (2005) on stock index options and Carr and Wu (2005a) on currency options

the two Lévy components, and the different correlations between the two Lévy components and their respective activity rates generate long-term skewness.

In contrast to modeling returns by a set of hidden factors, our modeling approach of applying stochastic time changes to different Lévy processes makes explicit the purpose of each modeling component. Under this framework, we use different Lévy processes as building blocks representing different economic forces. Applying stochastic time changes on each component randomizes the intensity of the impact from each economic force. The clear economic mapping makes the model design more intuitive and focused. Each component is added for a specific economic purpose. Using this approach is more likely to create models that are parsimonious and yet capable of delivering the required performance. For example, in Huang and Wu (2004), a diffusion return component captures small market movements and a jump component captures large market movements. Through different time changes, they ask whether stochastic volatility is driven by intensity changes of small or large movements or both, and whether the intensities on the two types of movements vary together or separately. In Carr and Wu (2005b), a diffusion component captures market movements at normal times and a jump to zero captures the impact of corporate default. Through time changes, the model captures the dynamic interactions between the market movement volatility and corporate default probability. In Carr and Wu (2004a), they use two Lévy components to capture separately the left and right tails of an exchange rate distribution, reflecting the impacts from the two economies. Through time changes, they randomize the activity rate underlying the two Lévy components and capture the tug of war between the two economies. Such explicit economic mappings greatly facilitates the interpretation and creation of new models capturing specific economic behaviors.

3.3. Theory and evidence on activity rate dynamics

Exploiting information in variance swap rates and various realized variance estimators constructed from high-frequency returns, Wu (2005) empirically study the activity rate dynamics for the S&P 500 index returns under a generalized affine framework. He finds that the activity rate for the index return contains an infinite-activity jump component, with its arrival rate proportional to the activity rate level. The Markov Chain Monte Carlo estimation in Eraker, Johannes, and Polson (2003) on long histories of index returns also suggest the presence of a jump component in the activity rate dynamics.

The impact of a jump component in the activity rate dynamics is usually small on the pricing of stock (index) options (Broadie, Chernov, and Johannes (2002)) and the term structure of variance swaps (Wu (2005)). Hence, many specifications for option pricing assume pure continuous activity rate dynamics for parsimony. Nevertheless, jumps are an integral part of the statistical variance dynamics. Furthermore, their pricing impacts can become more significant for derivative contracts that are sensitive to the tails of the variance distribution, e.g., options on variance swaps or realized variance.

When separate time changes are applied to different innovation components, the underlying activity rates can be modeled independently or with dynamic interactions. For example, Carr and Wu (2004a) assume that the two activity rates that govern the positive and negative Lévy components are independent of each other. Independent assumptions are also applied in the SV4 specification in Huang and Wu (2004). In contrast, Carr and Wu (2005b) find that stock return volatilities and corporate default arrival intensities co-move with each other. To capture the co-movements, they model the joint dynamics of the stock return diffusion variance rate v_t and the default arrival rate λ_t as,

$$\begin{aligned} dv_t &= (u_v - \kappa_v v_t) dt + \sigma_v \sqrt{v_t} dW_t^v, \\ \lambda_t &= \xi v_t + z_t, \\ dz_t &= (u_z - \kappa_z z_t - \kappa_{vz} v_t) dt + \sigma_z \sqrt{z_t} dW_t^z, \end{aligned} \tag{19}$$

where W_t^v and W_t^z denote two independent Brownian motions. The interactions between the diffusion variance and default arrival are captured by both the contemporaneous loading coefficient ξ and the dynamic predictive coefficient κ_{vz} .³

When the purpose is to capture the option price behavior at a narrow range of maturities, a one-factor activity rate specification is often adequate in generating stochastic volatilities. However, if the purpose is to capture the term structure of at-the-money implied volatilities or variance swap rates across a wide range of maturities, a one-factor activity rate process is often found inadequate. In most options markets, the persistence of the implied volatilities increases with the option maturity.

³Arguing that the default arrival may depend directly on the stock price level, Bob Goldstein further suggests an extended specification for equation (19) as $\lambda_t = \xi_v v_t + \xi_p \ln P_t + z_t$, where P_t denotes the stock price level.

This feature calls for multi-factor activity rate dynamics with different degrees of persistence for the different factors. One example is to allow the activity rate to revert to a stochastic mean level:

$$\begin{aligned} dv_t &= \kappa(m_t - \kappa_v v_t)dt + \sigma_v \sqrt{v_t} dW_t^v, \\ dm_t &= \kappa_m (\theta - m_t)dt + \sigma_m \sqrt{m_t} dW_t^m, \end{aligned}$$

where the mean-reversion speed of m , κ_m , is usually much smaller than the mean-reversion speed of the activity rate itself κ_v . Balduzzi, Das, and Foresi (1998) use a similar specification for the instantaneous interest rate dynamics and label m as the stochastic central tendency factor. Intuitively, the activity rate $v(t)$ affects short-term option implied volatilities more heavily whereas the central tendency factor m_t dominates the variation of long-term options. Thus, the persistence of the option implied volatility or variance swap rate can increase with the option maturities. Carr and Wu (2004a) consider a similar extension to their stochastic skew model, where the activity rates of both the positive and the negative Lévy components are allowed to revert to a common stochastic central tendency factor. Their estimation shows that the extension significantly improves the option pricing performance along the maturity dimension. Carr and Wu (2005a) also consider a similar extension on the default arrival dynamics to better capture the term structure of credit default swap spreads.

Most applications in option pricing use affine specifications for the activity rate dynamics, under which the activity rate is an affine function of a set of state variables and both the drift and variance of the state variables are affine in the state variable levels. When upward jumps are allowed in these state variables, their arrival rate are also affine in the state variables. Carr and Wu (2004b) show that both affine and quadratic specifications can be used to model the activity rate while retaining the analytical tractability for option pricing. Santa-Clara and Yan (2005) estimate a model with quadratic activity rates on S&P 500 index options. In their model, the return innovation consists of both a diffusion component and a compound Poisson jump component, and each component is time changed separately with the underlying activity rate being a quadratic function of an Ornstein-Uhlenbeck process.

Lewis (2000) and Heston (1997) show that option pricing is also reasonably tractable when the activity rate is governed by the 3/2 dynamics:

$$dv_t = \kappa v_t (\theta - v_t)dt + \sigma_v v_t^{3/2} dW_t. \quad (20)$$

Carr and Sun (2005) show that under a pure diffusion model under which the return variance rate follows this $3/2$ process, European option values can be written as a function of the asset price level and the level of the variance swap rate of the same maturity, with no separate dependence on calendar time or time-to-maturity. Furthermore, the pricing function depends only on the volatility of volatility coefficient σ_v , but not on the drift parameters (θ, κ) . Therefore, if we observe the underlying asset's price and its variance swap rate quotes, we can price options with merely one model parameter σ_v , without ever trying to estimate the drift function of the variance rate.

Within the one-factor diffusion context, several empirical studies find that a $3/2$ specification on the diffusion of the variance rate performs better the square-root specification. Favorable evidence based on time-series returns include Chacko and Viceira (2003), Ishida and Engle (2002), Javaheri (2005), and Jones (2003). Jones (2003) and Medvedev and Scaillet (2003) also find that the $3/2$ specification prices option better. Bakshi, Ju, and Ou-Yang (2004) find that the statistical dynamics of VIX^2 is also closer to a $3/2$ specification than to a square root specification.

4. Modeling financial security returns with time-changed Lévy processes

Once we have a clear understanding on the different roles played by Lévy innovations and random time changes, we can assemble the pieces together and write a complete model for the financial security return. The traditional literature often starts with the specification of the return dynamics under the statistical measure $\mathbb{P}$, and then derive the return dynamics under the risk-neutral measure $\mathbb{Q}$ for option pricing based on market price of risk specifications. However, since the requirement for analytical tractability mainly comes from the expectation operation under the risk-neutral measure in pricing contingent claims, it is often convenient to start directly with a tractable risk-neutral dynamics. Then, since we do not have as much concern for the tractability of the statistical dynamics, we can accommodate very flexible market price of risk specifications, with the only practical constraints coming from reasonability and identification considerations.

4.1. Constructing risk-neutral return dynamics

Let S_t denote the time- t price of a financial security. Let $\{X_{\tau_t^k}^k\}_{k=1}^K$ denote a series of independent time-changed Lévy processes, which are specified under a risk-neutral measure $\mathbb{Q}$. We use these processes as building blocks for the return dynamics. The independence assumption between different components is for convenience only, although interactions can be added when necessary as in equation (19). We model the risk-neutral return dynamics over the time period $[0, t]$ as,

$$\ln S_t / S_0 = (r - q)t + \sum_{k=1}^K \left(b^k X_{\tau_t^k}^k - \phi_{x^k}(b^k) \mathcal{T}_t^k \right), \quad (21)$$

where r denotes the instantaneous interest rate, q denotes the dividend yield for stocks or the foreign instantaneous interest rate for currencies, and b^k denotes a constant loading coefficient on the k th component. For notational clarity, I assume both r and q constant throughout the chapter. If we allow both to be stochastic, the first term should be replaced by an integral $\int_0^t (r(u) - q(u)) du$. If they vary deterministically over time, we can also replace the integral with the continuously compounded yields over the horizon $[0, t]$.

Equation (21) models the risks in the asset return using K components of time-changed Lévy processes. The cumulant exponent $\phi_{x^k}(b^k)$ represents a concavity adjustment so that the return dynamics satisfy the martingale condition under the risk-neutral measure:

$$\mathbb{E}_0^{\mathbb{Q}} [S_t / S_0] = e^{(r-q)t}. \quad (22)$$

Since by definition,

$$\mathbb{E}_0^{\mathbb{Q}} \left[e^{b^k X_t} \right] = e^{\phi_{x^k}(b^k)t}, \quad (23)$$

the following expectation is a martingale:

$$\mathbb{E}_0^{\mathbb{Q}} \left[e^{b^k X_t - \phi_{x^k}(b^k)t} \right] = 1. \quad (24)$$

The martingale condition retains when we replace t with a locally predictable and continuous time change $\mathcal{T}_t$ (Küchler and Sørensen (1997)):

$$\mathbb{E}_0^{\mathbb{Q}} \left[e^{b^k X_{\mathcal{T}_t^k} - \phi_{x^k}(b^k) \mathcal{T}_t^k} \right] = 1. \quad (25)$$

Thus, we have

$$\mathbb{E}_0^{\mathbb{Q}} [S_t/S_0] = \mathbb{E}_0^{\mathbb{Q}} \left[e^{(r-q)t + \sum_{k=1}^K \left(b^k X_{\mathcal{T}_t^k} - \phi_{x^k}(b^k) \mathcal{T}_t^k \right)} \right] = e^{(r-q)t} \prod_{k=1}^K \mathbb{E}_0^{\mathbb{Q}} \left[e^{b^k X_{\mathcal{T}_t^k} - \phi_{x^k}(b^k) \mathcal{T}_t^k} \right] = e^{(r-q)t}. \quad (26)$$

The independence assumption between different Lévy components enables us to move the expectation operation inside the product.

Each Lévy process X_t^k can have a drift component of its own, but it is irrelevant in our return specification (21) because any drift component will be canceled out with a corresponding term in the concavity adjustment. Hence, for each Lévy component, we only specify the diffusion volatility σ if the security price is allowed to move continuously and the Lévy density $\pi(x)$ if the price is allowed to jump.

The above return dynamics are defined over the horizon $[0, t]$ with time 0 referring to today and t being some future time corresponding to the maturity date of the contingent claim being valued. The time change $\mathcal{T}_t$ represents the integral of the activity rates over the same time period $[0, t]$. Sometimes it is more convenient to use t to denote the current date and T for future date, with $\tau = T - t$ denoting the time to maturity of the contingent claim. Then, the time change can be defined accordingly between this time period,

$$\mathcal{T}_{t,T} \equiv \int_t^T v_-(u) du. \quad (27)$$

The log return between $[t, T]$ can be written as,

$$\ln S_T/S_t = (r-q)(T-t) + \sum_{k=1}^K \left(b^k X_{\mathcal{T}_{t,T}^k} - \phi_{x^k}(b^k) \mathcal{T}_{t,T}^k \right). \quad (28)$$

4.1.1. Examples

We start with the simplest case where the return innovation is driven by one diffusion component without time change: $X_t^1 = \sigma W_t, \mathcal{T}_t = t, K = 1, b^1 = 1$, with W_t denoting a standard Brownian motion. The return process becomes,

$$\ln S_t/S_0 = (r - q)t + \sigma W_t - \frac{1}{2}\sigma^2 t. \quad (29)$$

The cumulant exponent of σW_t evaluated at $s = 1$ is $\phi(1) = \frac{1}{2}\sigma^2$. Equation (29) is essentially the classic Black and Scholes (1973) model.

Applying random time change to the diffusion component, we have,

$$\ln S_t/S_0 = (r - q)t + \sigma W_{\mathcal{T}_t} - \frac{1}{2}\sigma^2 \mathcal{T}_t, \quad (30)$$

where we simply replace the t with $\mathcal{T}_t$ on terms related to the Lévy component. If we model the activity rate underlying the time change v_t by the square-root process of Cox, Ingersoll, and Ross (1985), we will generate the stochastic volatility model of Heston (1993):

$$dv_t = \kappa(1 - v_t)dt + \sigma_v \sqrt{v_t} dW_t^v. \quad (31)$$

The long-run mean of the activity rate is normalized to one for identification purpose, since we already have a free volatility parameter σ in (30) that captures the mean level of volatility. In the original Heston model, σ is normalized to one and the long-run mean of the activity rate is left as a free parameter. To match the activity rate specification with the time change notation, we can rewrite the activity rate in integral forms,

$$v_t = v_0 + \int_0^t \kappa(1 - v_t) dt + \int_0^t \sigma_v \sqrt{v_t} dW_t^v = v_0 + \kappa t - \kappa \mathcal{T}_t + \sigma_v W_{\mathcal{T}_t}^v. \quad (32)$$

Technically, the second equality in (32) holds only *in distribution* and the W^v in $\int_0^t \sqrt{v_t} dW_t^v$ denotes a different Brownian motion from the W^v in $W_{\mathcal{T}_t}^v$. Hence, a more technically correct way of writing the equality is:

$$\int_0^t \sqrt{v_t} dW_t^v =^d \tilde{W}_{\mathcal{T}_t}^v, \quad (33)$$

where $=^d$ denotes “equality in distribution,” and $(W^v, \tilde{W}^v)$ denote two different (and independent) Brownian motions. To avoid clustering, we abuse the notation in this chapter and use W^v to represent two different Brownian motions in the two different representations. We also use the same equality sign “=” to represent both the traditional mathematical equality and the equality in distribution. Analogously, the equalities between $\int_0^t \sqrt{v_t} dW_t$ and $W_{\mathcal{T}_t}$, between $\sqrt{v_t} dW_t$ and $dW_{\mathcal{T}_t}$, and between $\sqrt{v_t} dW_t^v$ and $dW_{\mathcal{T}_t}^v$ are all in distribution, and the two “ W ”s in each pair represent two different Brownian motions. Heston (1993) allows correlation between the activity rate innovation and the return innovation $\mathbb{E}[dW_t dW_t^v] = \rho dt$, or equivalently under the time-change notation, $\mathbb{E}[dW_{\mathcal{T}_t} dW_{\mathcal{T}_t}^v] = \rho d\mathcal{T}_t = \rho v_t dt$.

Technicality aside, I regard the time-change notation as simply a convenient way of rewriting the traditional stochastic differential equation. Using the Heston (1993) model as an example. The traditional representation in terms of the stochastic differential equation is:

$$\begin{aligned} d \ln S_t &= (r - q)dt + \sigma \sqrt{v_t} dW_t - \frac{1}{2} \sigma^2 v_t dt, \\ dv_t &= \kappa(1 - v_t)dt + \sigma_v \sqrt{v_t} dW_t^v. \end{aligned} \quad (34)$$

The following time-changed Lévy process generates the same return distribution:

$$\begin{aligned} \ln S_t / S_0 &= (r - q)t + \sigma W_{\mathcal{T}_t} - \frac{1}{2} \sigma^2 \mathcal{T}_t, \\ v_t &= v_0 + \kappa t - \kappa \mathcal{T}_t + \sigma_v W_{\mathcal{T}_t}^v, \end{aligned} \quad (35)$$

with the technical caveat that (34) and (35) represent different processes and (W, W^v) in the two sets of equations represent completely different Brownian motions.

Now consider an example where the return innovation is driven by a pure jump Lévy component without time change and the jump arrival is governed by the dampened power law specification in (12) with $\alpha \neq 0$ and $\alpha \neq 1$. The return dynamics can be written as,

$$\ln S_t / S_0 = (r - q)t + J_t - \phi_J(1)t, \quad (36)$$

where J_t denotes this Lévy jump component, and the cumulant exponent is,

$$\phi_J(s) = \Gamma(-\alpha) \lambda [(\beta_+ - s)^\alpha - \beta_+^\alpha + (\beta_- + s)^\alpha - \beta_-^\alpha] + sC(h), \quad (37)$$

with $C(h)$ given in (14). Since any linear drift terms in J_t will be canceled out by the corresponding term in the concavity adjustment, it becomes obvious that the exact form of the truncation function and the resultant linear coefficient $C(h)$ are immaterial for modeling and estimation. Given the cumulant exponent in (13), the concavity adjustment in equation (36) becomes,

$$\phi_J(1) = \Gamma(-\alpha)\lambda [(\beta_+ - 1)^\alpha - \beta_+^\alpha + (\beta_- + 1)^\alpha - \beta_-^\alpha] + C(h). \quad (38)$$

If we apply random time change to the pure jump Lévy component, we can simply replace J_t with $J_{\mathcal{T}_t}$ and $\phi_J(1)t$ with $\phi_J(1)\mathcal{T}_t$:

$$\ln S_t/S_0 = (r - q)t + J_{\mathcal{T}_t} - \phi_J(1)\mathcal{T}_t, \quad (39)$$

which is a pure jump process with stochastic volatility generated purely from the stochastic arrival of jumps.

When we use one Lévy component in the return dynamics, it is natural to set the loading coefficient b to unity as it can always be absorbed into the scaling specification of the Lévy process. To show an example where the loading coefficient plays a more explicit role, we consider a market model for stock returns, where the return on each stock is decomposed into two orthogonal components: a market risk component and an idiosyncratic risk component. We use a Lévy process X_t^m to model the market risk and another Lévy process X_t^j to model the idiosyncratic risk for stock j . Then, the return on stock j can be written under the risk-neutral measure $\mathbb{Q}$ as,

$$\ln S_t^j/S_0^j = (r - q)t + (b^j X_t^m - \phi_{X^m}(b^j)t) + (X_t^j - \phi_{X^j}(1)t), \quad (40)$$

where the first component $(r - q)t$ captures the instantaneous drift under the risk-neutral measure, the second component $(b^j X_t^m - \phi_{X^m}(b^j)t)$ represents the concavity-adjusted market risk component, with b^j capturing the linear loading of the return on the market risk factor X_t^m , and the last component $(X_t^j - \phi_{X^j}(1)t)$ is the concavity-adjusted idiosyncratic risk component for the stock return.

Under the Lévy specification in (40), stock returns are iid. We can apply random time changes to the two Lévy processes to generate stochastic volatility:

$$\ln S_t^j / S_0^j = (r - q)t + \left(b^j X_{\tau_t^m}^m - \phi_{x^m}(b^j) \tau_t^m \right) + \left(X_{\tau_t^j}^j - \phi_{x^j}(1) \tau_t^j \right), \quad (41)$$

where stochastic volatility can come either from the market risk via τ_t^m or from the idiosyncratic risk via τ_t^j . Mo and Wu (2005) propose an international capital asset pricing model with a similar structure, where X^m represents a global risk factor and X^j a country-specific risk factor. They specify the dynamics under both the risk-neutral measure and the statistical measure, and estimate the joint dynamics of three economies (US, UK, and Japan) using the time-series returns and option prices on the S&P 500 index, the FTSE 100 Index, and the Nikkei-225 Stock Average.

4.2. Market price of risks and statistical dynamics

Once we have specified the return dynamics under the risk-neutral measure $\mathbb{Q}$, we can derive the dynamics under the statistical measure $\mathbb{P}$ if we know whether and how different sources of risks are priced. Take the generic return specification in (21) as an example. We have K sources of return risks as captured by the K Lévy process components $\{X_t^k\}_{k=1}^K$. We also have K sources of volatility risks corresponding to each return component. Furthermore, each Lévy process X_t^k can have a diffusion component and a jump component. The two components can be priced differently. Upside and downside jumps can also be priced differently (Wu (2006) and Bakshi and Wu (2005)). The activity rate that underlies each time change can also have a diffusion and a jump component that can be priced differently. Depending on the market price of risk specification, the statistical return dynamics can look dramatically different from the risk-neutral dynamics.

In this subsection, we consider a simple class of market price of risk specifications, which in most cases generates statistical return dynamics that stay in the same class as the risk-neutral dynamics. The pricing kernel that defines the market price of all sources of risks can be written as,

$$\mathcal{M}_t = e^{-rt} \prod_{k=1}^K \exp \left(-\gamma_k X_{\tau_t^k}^k - \phi_{x^k}(-\gamma_k) - \gamma_{kv} X_{\tau_t^k}^{kv} - \phi_{x^{kv}}(-\gamma_{kv}) \right) \cdot \zeta, \quad (42)$$

where $X_{\mathcal{T}_t^k}^k$ denotes the return risk as in (21), $X_{\mathcal{T}_t^k}^{kv}$ denotes another set of time-changed Lévy processes that characterize the activity rate risk, and ζ denotes an orthogonal martingale component that prices other sources of risks independent of the security return under consideration. We maintain the constant interest rate assumption in the pricing kernel specification. The exponential martingale components in the pricing kernel determines the measure change from $\mathbb{P}$ to $\mathbb{Q}$. The simplicity of the specification comes from the constant assumption on the market price coefficients γ_k and γ_{kv} .

Given the pricing kernel specification in (42) and the risk-neutral return dynamics in (21), we can infer the statistical return dynamics. We use examples to illustrate the procedure, starting with the simplest case where the return is driven by one diffusion component without time change as in (29). According to the above exponential martingale assumption, the measure change from $\mathbb{P}$ to $\mathbb{Q}$ is defined by,

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_t = \exp \left(-\gamma \sigma W_t - \frac{1}{2} \gamma^2 \sigma^2 t \right). \quad (43)$$

with $\phi_{\sigma W}(-\gamma) = \frac{1}{2} \gamma^2 \sigma^2$. The literature has taken different approaches in arriving at the dynamics under a measure change. For measure changes defined by exponential martingales of a Lévy processes X , it is convenient to remember that $\phi_X^{\mathbb{P}}(s) = \phi_X^{\mathbb{Q}}(s + \gamma) - \phi_X^{\mathbb{Q}}(\gamma)$ and that the drift adjustment of X is captured by $\eta = \phi^{\mathbb{P}}(1) - \phi^{\mathbb{Q}}(1)$ (Küchler and Sørensen (1997)). For the simple case in (43) with $X = \sigma W$, we have $\phi_{\sigma W}^{\mathbb{Q}}(1) = \frac{1}{2} \sigma^2$, and

$$\phi_{\sigma W}^{\mathbb{P}}(1) = \phi_{\sigma W}^{\mathbb{Q}}(1 + \gamma) - \phi_{\sigma W}^{\mathbb{Q}}(\gamma) = \frac{1}{2} (1 + \gamma)^2 \sigma^2 - \frac{1}{2} \gamma^2 \sigma^2 = \frac{1}{2} \sigma^2 + \gamma \sigma^2. \quad (44)$$

Thus, the drift adjustment, or the instantaneous expected excess return, is $\eta = \gamma \sigma^2$. The statistical ($\mathbb{P}$) return dynamics becomes,

$$\ln S_t / S_0 = (r - q)t + \gamma \sigma^2 t + \sigma W_t - \frac{1}{2} \sigma^2 t. \quad (45)$$

For the Heston (1993) model, which has the risk-neutral dynamics specified in (34) or equivalently (35), the associated exponential martingale that defines the measure change from $\mathbb{P}$ and $\mathbb{Q}$ becomes,

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_t = \exp \left(-\gamma \sigma W_{\mathcal{T}_t} - \frac{1}{2} \gamma^2 \sigma^2 \mathcal{T}_t - \gamma_v \sigma_v W_{\mathcal{T}_t}^v - \frac{1}{2} \gamma_v^2 \sigma_v^2 \mathcal{T}_t \right). \quad (46)$$

The cumulant exponent of the return innovation σW_t under measure $\mathbb{P}$ becomes $\phi_{\sigma W}^{\mathbb{P}}(s) = \phi_{\sigma W}^{\mathbb{Q}}(s + \gamma + \gamma_v \sigma_v \rho / \sigma) - \phi_{\sigma W}^{\mathbb{Q}}(\gamma + \gamma_v \sigma_v \rho / \sigma)$. Hence, the drift adjustment induced by the measure change is $\eta = \gamma \sigma^2 + \gamma_v \sigma_v \sigma \rho$. The first term is induced by the pricing of the return risk W and the second term is induced by the pricing of the part of volatility risk W^v that is correlated with the return risk. Given the stochastic time change and hence stochastic activity rate, the risk premium over the horizon $[0, t]$ is $\eta \mathcal{T}_t$, and the instantaneous risk premium at time t is ηv_t . The statistical return dynamics becomes,

$$\ln S_t / S_0 = (r - q)t + (\gamma \sigma^2 + \gamma_v \sigma_v \sigma \rho) \mathcal{T}_t + \sigma W_{\mathcal{T}_t} - \frac{1}{2} \sigma^2 \mathcal{T}_t. \quad (47)$$

To derive the statistical dynamics for the activity rate, we note that the cumulant exponent of the activity rate innovation $\sigma_v W^v$ under measure $\mathbb{P}$ becomes $\phi_{\sigma_v W^v}^{\mathbb{P}}(s) = \phi_{\sigma_v W^v}^{\mathbb{Q}}(s + \gamma_v + \gamma \sigma \rho / \sigma_v) - \phi_{\sigma_v W^v}^{\mathbb{Q}}(\gamma_v + \gamma \sigma \rho / \sigma_v)$. Hence, the measure change induces an instantaneous drift change captured by $\eta^v = \phi_{\sigma_v W^v}^{\mathbb{P}}(1) - \phi_{\sigma_v W^v}^{\mathbb{Q}}(1) = \gamma_v \sigma_v^2 + \gamma \sigma \sigma_v \rho$, where the first term is induced by the pricing of the activity rate innovation W^v and the second term is induced by the pricing of the part of return risk W that is correlated with the activity rate. Since we apply the same time change $\mathcal{T}_t$ to the two sources of risks W and W^v , the actual drift adjustment over calendar time $[0, t]$ becomes $\eta^v \mathcal{T}_t$, and the instantaneous adjustment is $\eta^v v_t$. The statistical activity rate dynamics becomes,

$$v_t = v_0 + at - \kappa \mathcal{T}_t + \eta^v \mathcal{T}_t + \sigma_v W_{\mathcal{T}_t}^v, \quad (48)$$

or in the form of the stochastic differential equation,

$$dv_t = (a - (\kappa - \eta^v) v_t) dt + \sigma_v \sqrt{v_t} dW_t^v, \quad (49)$$

where the measure change induces a change in the mean reversion speed from κ under $\mathbb{Q}$ to $\kappa^{\mathbb{P}} = \kappa - \eta^v = \kappa - \gamma_v \sigma_v^2 - \gamma \sigma \sigma_v \rho$ under $\mathbb{P}$. Estimation on stock indexes and stock index options often find that the market price of return risk (γ) is positive and the market price of variance risk (γ_v) is negative. Given the well-documented negative correlation (ρ) between the return and variance innovations, both

sources of market prices make the activity rate more persistent under the risk-neutral measure than the activity rate is under the statistical measure: $\kappa < \kappa^{\mathbb{P}}$.⁴

For the pure jump Lévy process example as in (36), the measure change from $\mathbb{P}$ to $\mathbb{Q}$ is defined by the exponential martingale,

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_t = \exp(-\gamma J_t - \phi_J(-\gamma)t). \quad (50)$$

The Lévy density under the two measures are linked by $\pi^{\mathbb{P}}(x) = e^{\gamma x} \pi^{\mathbb{Q}}(x)$. If the Lévy density under $\mathbb{Q}$ is given by equation (12), its corresponding Lévy density under $\mathbb{P}$ becomes,

$$\pi^{\mathbb{P}}(x) = \begin{cases} \lambda \beta_+ \exp(-(\beta_+ - \gamma)x) x^{-\alpha-1}, & x > 0, \\ \lambda \beta_- \exp(-(\beta_- + \gamma)|x|) |x|^{-\alpha-1}, & x < 0, \end{cases} \quad (51)$$

Therefore, the Lévy density is still controlled by a dampened power law under the statistical measure $\mathbb{P}$, only with the exponential dampening coefficients changed from (β_+, β_-) under $\mathbb{Q}$ to $\beta_+^{\mathbb{P}} = \beta_+ - \gamma$ and $\beta_-^{\mathbb{P}} = \beta_- + \gamma$ under $\mathbb{P}$. The dampening coefficients should be nonnegative under both measures. This condition limits the range of values that the market price of risk γ can take. Given the risk-neutral dampening coefficients (β_+, β_-) , we need $\gamma \in [-\beta_-, \beta_+]$. Given the statistical coefficients $(\beta_+^{\mathbb{P}}, \beta_-^{\mathbb{P}})$, we need $\gamma \in [-\beta_+^{\mathbb{P}}, \beta_-^{\mathbb{P}}]$.

Wu (2006) and Bakshi and Wu (2005) allow the downside and upside jumps to have different market prices (γ_+, γ_-) . In this case, we can directly specify the dampening coefficients under the two measures (β_+, β_-) and $(\beta_+^{\mathbb{P}}, \beta_-^{\mathbb{P}})$ as free parameters with positivity constraints. Then, the market prices of positive and negative jump risks can be derived as $\gamma_+ = \beta_+ - \beta_+^{\mathbb{P}}$ and $\gamma_- = \beta_-^{\mathbb{P}} - \beta_-$. By estimating this pure jump Lévy model to S&P 500 index time-series returns and option prices, Wu finds that there is zero dampening on downside jumps under the risk-neutral measure ($\beta_- = 0$). Thus, the market price of downside jump risk reaches its upper limit at $\gamma_- = \beta_-^{\mathbb{P}}$. This extremely high market price of downside risk is needed to capture the much higher prices for out-of-the-money put options than for the corresponding out-of-the-money call options on the index and the corresponding implied volatility smirk at both short and long maturities.

⁴Since the constant part of drift remains the same as a , the long-run mean of the activity rate changes from a/κ under $\mathbb{Q}$ to $a/(\kappa - \eta^v)$ under $\mathbb{P}$. The smaller mean reversion under $\mathbb{Q}$ implies a higher long-run mean.

Given the measure change defined in (50), the cumulant exponent under measure $\mathbb{P}$ is linked to the cumulant exponent under measure $\mathbb{Q}$ by $\varphi^{\mathbb{P}}(s) = \varphi^{\mathbb{Q}}(s + \gamma) - \varphi^{\mathbb{Q}}(\gamma)$. The instantaneous expected excess return is given by $\eta = \varphi_J^{\mathbb{P}}(1) - \varphi_J^{\mathbb{Q}}(1) = \varphi_J^{\mathbb{Q}}(1 + \gamma) - \varphi_J^{\mathbb{Q}}(\gamma) - \varphi_J^{\mathbb{Q}}(1)$. It is obvious that any term in the cumulant exponent $\varphi_J^{\mathbb{Q}}(s)$ that is linear in s does not contribute to the expected excess return η . Hence, the truncation-induced linear term $sC(h)$, or the choice of the truncation function $h(x)$, does not affect the computation of the expected excess return η .

Under the jump specification in (12) and when $\alpha \neq 0$ and $\alpha \neq 1$, the instantaneous expected excess return is:

$$\begin{aligned} \eta = & \Gamma(-\alpha)\lambda [((\beta_+ - \gamma) - 1)^\alpha - (\beta_+ - \gamma)^\alpha + ((\beta_- + \gamma) + 1)^\alpha - (\beta_- + \gamma)^\alpha] \\ & - \Gamma(-\alpha)\lambda [(\beta_+ - 1)^\alpha - \beta_+^\alpha + (\beta_- + 1)^\alpha - \beta_-^\alpha], \end{aligned} \quad (52)$$

where the first line is the cumulant exponent under measure $\mathbb{P}$ evaluated at $s = 1$ and the second line is the cumulant exponent under measure $\mathbb{Q}$ evaluated at $s = 1$, with the term $C(h)$ in both cumulant exponents dropping out. Nevertheless, sometimes the measure change itself can induce an additional linear term that contributes to the expected excess return. Hence, it is safer to always evaluate η according to the equation $\eta = \varphi_J^{\mathbb{Q}}(1 + \gamma) - \varphi_J^{\mathbb{Q}}(\gamma) - \varphi_J^{\mathbb{Q}}(1)$.

If we apply random time changes to the Lévy jump process and if the underlying activity rate risk is not correlated with the Lévy jump risk, we can simply replace ηt with $\eta \mathcal{T}_t$ as the excess return over time period $[0, t]$. If the risk-neutral activity rate follows a square root process, the statistical dynamics for the activity rate can be derived analogous to (49) with $\rho = 0$ due to the orthogonality between the jump innovation in return and the diffusion innovation in the activity rate.

4.3. More flexible market price of risk specifications

Under the exponential martingale specification embedded in the pricing kernel in (42), return and volatility risks are both captured by a vector of time-changed Lévy processes, $(X_{\mathcal{T}_t}^k, X_{\mathcal{T}_t}^{kv})_{k=1}^K$ and the market prices on these risks, $(\gamma_k, \gamma_{kv})_{k=1}^K$ are assumed to be constant. The specification is parsimonious, under which the return (and activity rate) dynamics often stay within the same class under the two

measures $\mathbb{P}$ and $\mathbb{Q}$. However, since tractability requirement mainly comes from option pricing due to the associated expectation operation under the risk-neutral measure, a more flexible market price of risk specification poses little problems if we start the modeling with a tractable risk-neutral dynamics. Complex market price of specifications only lead to complex statistical dynamics, which are irrelevant for option pricing. The complication does affect the derivation of the likelihood functions for time-series estimation. Yet, when the return series can be sampled frequently, an euler approximation of the statistical dynamics often works well for estimation and it avoids the complication of taking expectations under the statistical measure for the conditional density derivation. Hence, we can freely specify arbitrarily complex market price of risks without incurring any difficulty for asset pricing. Beside the usual technical conditions that a pricing kernel needs to satisfy, the only practical constraints for the market price of risk specification come from reasonability and identification considerations. Even if a specification is mathematically allowed, we may discard it if it does not make economic sense and does not represent common investor behavior. Furthermore, a more flexible specification on the market price of risk gives us more degree of freedom, but it can also cause difficulties in identification. Hence, it is always prudent to start with a parsimonious assumption on the market price of risk and consider extensions only when the data ask for it.

Take the Black-Scholes model as a simple example, where the stock return under the risk-neutral measure $\mathbb{Q}$ is normally distributed with constant volatility σ as described in (29). Now we consider a super flexible market price of risk specification that defines the measure change from $\mathbb{P}$ to $\mathbb{Q}$ as,

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_t = \exp \left(-(\gamma_0 + \gamma_1 Z_t + \gamma_2 Z_t^2 + \gamma_3 Z_t^3) \sigma W_t - \frac{1}{2} (\gamma_0 + \gamma_1 Z_t + \gamma_2 Z_t^2 + \gamma_3 Z_t^3)^2 \sigma^2 t \right), \quad (53)$$

where the market price of risk is given by a polynomial function of Z_t , $\gamma = \gamma_0 + \gamma_1 Z_t + \gamma_2 Z_t^2 + \gamma_3 Z_t^3$ with Z_t being some state variable whose dynamics is left unspecified. The order of three is purely arbitrary and for illustration only. Then, the instantaneous expected excess return at time t is $\eta_t = (\gamma_0 + \gamma_1 Z_t + \gamma_2 Z_t^2 + \gamma_3 Z_t^3) \sigma^2$, and the security price becomes,

$$dS_t/S_t = (r - q + (\gamma_0 + \gamma_1 Z_t + \gamma_2 Z_t^2 + \gamma_3 Z_t^3) \sigma^2) dt + \sigma dW_t. \quad (54)$$

Many empirical studies identify dividend yield, default spread, interest rate, and so on as variables that can predict expected excess returns. If the evidence is robust, we can use them as the state variable Z_t , which can either be a scalar or a vector.

Regardless of the complexity of the statistical dynamics, option pricing still follows the Black-Scholes formula. The return distribution under the statistical measure $\mathbb{P}$ depends on the dynamics of Z_t . Nevertheless, with an euler approximation, we can still assume that the conditional return distribution over a short time interval $[t, t + \Delta t]$ is normally distributed, with mean $(r - q + \eta_t - \frac{1}{2}\sigma^2)\Delta t$ and variance $\sigma^2\Delta t$, and then construct the conditional likelihood function of the return accordingly.

Consider another example where the activity rate follows a square-root dynamics under the risk-neutral measure,

$$dv_t = (a - \kappa v_t)dt + \sigma_v \sqrt{v_t} dW_t^v. \quad (55)$$

For simplicity, we assume that the return Lévy innovation is not correlated with the Brownian motion W_t^v in the activity rate process. As shown in a later section, the affine structure of the activity rate dynamics under the risk-neutral measure makes option pricing tractable. The previous section assumes a constant market price γ_v on $\sigma_v \sqrt{v_t} dW_t^v$, which induces a drift change of $\gamma_v \sigma_v^2 v_t$. Hence, it amounts to change the mean-reversion coefficient from κ to $\kappa - \gamma_v \sigma_v^2$ under $\mathbb{P}$. Now we consider a more general specification,

$$\gamma_t^v = \gamma_0/v_t + \gamma_1 + \gamma_2 v_t + \cdots + \gamma_k v_t^{k-1}, \quad (56)$$

for any order k . The induced drift change becomes: $\gamma_0 \sigma_v^2 + \gamma_1 \sigma_v^2 v_t + \gamma_2 \sigma_v^2 v_t^2 + \cdots + \gamma_k \sigma_v^2 v_t^k$. Thus, the drift of the activity rate process is no longer affine under the statistical measure, but the complication does not affect option pricing and we can resort to euler approximation for the likelihood construction.

Nevertheless, the specification in (56) is not completely innocuous. When v_t approaches zero, its risk (innovation) $\sigma_v \sqrt{v_t} dW_t^v$ also approaches zero, yet the risk premium does not approach zero, but approaches a non-zero constant $\gamma_0 \sigma_v$. A riskless security cannot earn a non-zero risk premium. Hence, the specification violates the no-arbitrage condition if the activity rate can stay at zero. Recently, Cheridito, Filipović, and Kimmel (2003) and Pan and Singleton (2005) apply a restricted version of (56) with $\gamma_k = 0$ for $k \geq 2$. Then, the risk premium is affine in v_t , and hence the activity rate dynamics remain affine under the statistical measure. To guarantee no-arbitrage, they add further technical conditions on

the statistical dynamics so that zero is not an absorbing barrier, but a reflecting barrier of v_t , and hence v_t does not stay at zero long enough for investors to do arbitrage trading. The technical condition guarantees no arbitrage. Nevertheless, the specification with γ_0 strictly nonzero still implies that investors charge a risk premium no smaller than $\gamma_0 \sigma_v^2$ no matter how small the risk becomes. A flexible market price of risk specification does not hinder option pricing as long as we start with the risk-neutral dynamics, but it remains important to apply our economic sense and the rule of parsimony and discipline in specifying them.

In the fixed income literature, an enormous amount of studies exploit various forms of the expectation hypothesis to predict future exchange rate movements using current interest rate differentials between the two economies, and predict short-term interest rate movements with the current term structure information. Several recent studies explore whether affine models can explain the regression slope coefficients.⁵ Affine models ask that bond yields of all maturities are affine functions of a set of state variable. This cross-sectional relation has bearings on the risk-neutral dynamics: The risk-neutral drift and variance of the state vector are both affine functions of the state vector. However, it has no direct bearings on the statistical dynamics. Hence, we can retain the affine cross-sectional relation between bond yields of different maturities while allowing flexible specifications for the market price of risks and hence flexible drift functions for the state vector under the statistical measure. Nevertheless, most studies impose discipline by requiring that the statistical drift of the state vector is also affine. This self-imposed requirement limits the market price of risk specification to an affine form $\gamma(X_t) = a + bX_t$ when the state variable has a constant diffusion, and of the form $\gamma(X_t) = a/X_t + b$ as in (56) with $\gamma_k = 0$ for $k \geq 2$ when the state variable follows a square root process.

⁵Examples include Backus, Foresi, Mozumdar, and Wu (2001), Duffee (2002), Dai and Singleton (2002), and Roberds and Whiteman (1999) for expectation hypotheses on the term structure in a single economy, and Backus, Foresi, and Telmer (2001), for international term structure and currency pricing.

5. Option pricing under time-changed Lévy processes

To price options when the underlying asset return is driven by Lévy processes with or without time changes, we first derive the generalized Fourier transform of the asset return under the risk-neutral measure and then use a Fourier inversion method to numerically compute option prices.

5.1. Deriving the Fourier transforms

Carr and Wu (2004b) propose a theorem that significantly improves the tractability of option pricing under time-changed Lévy processes. They convert the problem of finding the generalized transform of a time-changed Lévy process into the problem of finding the Laplace transform of the random time under a new complex-value measure,

$$\phi_Y(u) \equiv \mathbb{E}^{\mathbb{Q}} [e^{iuX_{T_t}}] = \mathbb{E}^{\mathbb{M}} [e^{-\psi_x(u)T_t}], \quad (57)$$

where $\psi_x(u)$ denotes the characteristic exponent of the underlying Lévy process X_t , and the second expectation is under a new measure $\mathbb{M}$, defined by the following complex-valued exponential martingale:

$$\left. \frac{d\mathbb{M}}{d\mathbb{Q}} \right|_t = \exp(iuX_{T_t} + T_t\psi_x(u)). \quad (58)$$

When the activity rate ν_t underlying the time change is independent of the Lévy innovation X_t , the measure change is not necessary and the result in (57) can be obtained via the law of iterated expectations. When the two processes are correlated, the proposed measure change simplifies the calculation by absorbing the effect of correlation into the new measure.

According to (57), tractable Fourier transforms for the time-changed Lévy process, $\phi_Y(u)$, can be obtained if we can obtain tractable forms for the characteristic exponent of the Lévy process, $\psi_x(u)$, and the Laplace transform of the time change. The three most widely used Lévy jump specifications include the Merton (1976) compound Poisson models with normally distributed jump sizes, the dampened power law specification and its various special cases, and the normal inverse gamma model and its extensions. All these models have analytical solutions for the characteristic exponents.

To solve for the Laplace transform, it is important to note that if we write the time change $\mathcal{T}_t$ in terms of the activity rate $\mathcal{T}_t = \int_0^t v(s_-)ds$, the same form of expectation appears in the bond pricing literature with the analogous term for the instantaneous activity rate being the instantaneous interest rate. Furthermore, since both nominal interest rates and the activity rate are required to be positive, they can be modeled using similar dynamics. Therefore, any interest rate dynamics that generate tractable bond pricing formulas can be borrowed to model the activity rate dynamics under measure $\mathbb{M}$ with tractable solutions to the Laplace transform in equation (57). In particular, the affine class of Duffie and Kan (1996), Duffie, Pan, and Singleton (2000), Duffie, Filipović, and Schachermayer (2003) and the quadratic class of Leippold and Wu (2002) for interest rates can be borrowed to model the activity rate dynamics with tractable exponential affine and exponential quadratic solutions for the Laplace transform, respectively. Carr and Wu (2004b) discuss these models in their general forms. Of all these specifications, the most popular is the square root process used in Heston (1993) and its various extensions to multiple factors and to include positive jumps. The 3/2 specification also generates tractable solutions for the Laplace transform in (57), but the solution contains a confluent hypergeometric function $M(\alpha, \beta; z)$, where the two coefficients (α, β) are complex valued and are functions of the characteristic coefficient u , and the argument z is a function of the activity rate level and option maturity. It remains a numerical challenge to compute this function efficiently over the wide range of complex-valued coefficients necessary for option pricing.

I illustrate the valuation procedure using the simple examples discussed in the previous sections, starting with the Black-Scholes model with the risk-neutral return dynamics given in (29):

$$\phi_s(u) \equiv \mathbb{E}^{\mathbb{Q}} \left[e^{iu \ln S_t / S_0} \right] = e^{iu(r-q)t} \mathbb{E}^{\mathbb{Q}} \left[e^{iu(\sigma W_t - \frac{1}{2}\sigma^2 t)} \right] = e^{iu(r-q)t - \frac{1}{2}(iu + u^2)\sigma^2 t}. \quad (59)$$

Given the constant interest rate and dividend yield assumption, we can factor them out before taking the expectation. In this case, the concavity adjustment term $iu\frac{1}{2}\sigma^2 t$ can also be factored out. Nevertheless, with time changes in my mind, I leave it inside the expectation and write $\psi_x(u) = \frac{1}{2}(iu + u^2)\sigma^2$ as the characteristic exponent of the *concavity-adjusted* return innovation term: $X_t = \sigma W - \frac{1}{2}\sigma^2 t$.

The Black-Scholes option pricing formula is well-known, deriving the generalized Fourier transform under the Black-Scholes model merely serves as a benchmark for more complicated examples. The first extension is to apply random time changes to the Black-Scholes specification,

$$\ln S_t/S_0 = (r - q)t + \sigma W_{\tau_t} - \frac{1}{2}\sigma^2 \tau_t. \quad (60)$$

Here, we can apply Carr and Wu's theorem to find the generalized Fourier transform:

$$\phi_s(u) = e^{iu(r-q)t} \mathbb{E}^{\mathbb{Q}} \left[e^{iu(\sigma W_{\tau_t} - \frac{1}{2}\sigma^2 \tau_t)} \right] = e^{iu(r-q)t} \mathbb{E}^{\mathbb{M}} \left[e^{-\psi_x(u)\tau_t} \right], \quad (61)$$

with $\psi_x(u) = \frac{1}{2}(iu + u^2)\sigma^2$ the same as for the concavity-adjusted return innovation for the Black-Scholes model. The construction of the new measure $\mathbb{M}$ and the Laplace transform under this new measure depend on the specification of the activity rate dynamics.

Take the Heston (1993) model as an example, where the activity rate dynamics under measure $\mathbb{Q}$ is, in stochastic differential equation form,

$$dv_t = \kappa(1 - v_t)dt + \sigma_v \sqrt{v_t} dW_t^v, \quad \rho dt = \mathbb{E} [dW_t dW_t^v]. \quad (62)$$

The measure change is defined by

$$\left. \frac{d\mathbb{M}}{d\mathbb{Q}} \right|_t = \exp \left(iu \left(\sigma W_{\tau_t} - \frac{1}{2}\sigma^2 \tau_t \right) + \tau_t \psi_x(u) \right). \quad (63)$$

The probabilistically equivalent writing under more traditional notation is,

$$\left. \frac{d\mathbb{M}}{d\mathbb{Q}} \right|_t = \exp \left(iu\sigma \int_0^t \sqrt{v_s} dW_s + \frac{1}{2}u^2\sigma^2 \int_0^t v_s ds \right), \quad (64)$$

where I plug in $\psi_x(u)$ and cancel out the concavity-adjustment term. This measure change induces a drift change in the activity rate dynamics given by the covariance term:⁶

$$\mu(v)^{\mathbb{M}}dt - \mu(v)^{\mathbb{Q}}dt = \langle iu\sigma\sqrt{v_t}dW_t, \sigma_v\sqrt{v_t}dW_t^v \rangle = iu\sigma\sigma_v v_t \rho dt. \quad (66)$$

Hence, under measure $\mathbb{M}$, the activity rate dynamics become,

$$dv_t = \left(\kappa - \kappa^{\mathbb{M}} v_t \right) dt + \sigma_v \sqrt{v_t} dW_t^v, \quad \kappa^{\mathbb{M}} = \kappa - iu\sigma\sigma_v \rho. \quad (67)$$

Both the drift and the instantaneous variance are affine in v_t under measure $\mathbb{M}$. Hence, the Laplace transform in (61) is exponential affine in the current level of the activity rate:

$$\phi_s(u) = e^{iu(r-q)t} \mathbb{E}^{\mathbb{M}} \left[e^{-\psi_x(u)T_t} \right] = e^{iu(r-q)t - b(t)v_0 - c(t)}, \quad (68)$$

with the coefficients $b(t)$ and $c(t)$ given by,

$$\begin{aligned} b(t) &= \frac{2\psi_x(u)(1-e^{-\xi t})}{2\xi - (\xi - \kappa^{\mathbb{M}})(1-e^{-\xi t})}, \\ c(t) &= \frac{\kappa}{\sigma_v^2} \left[2 \ln \left(1 - \frac{\xi - \kappa^{\mathbb{M}}}{2\xi} (1 - e^{-\xi t}) \right) + (\xi - \kappa^{\mathbb{M}})t \right], \end{aligned} \quad (69)$$

with $\xi = \sqrt{(\kappa^{\mathbb{M}})^2 + 2\sigma_v^2\psi_x(u)}$.

Suppose we further allow the activity rate to revert to a stochastic central tendency factor in generating a two-factor activity rate dynamics under measure $\mathbb{Q}$:

$$\begin{aligned} dv_t &= \kappa(m_t - v_t)dt + \sigma_v \sqrt{v_t} dW_t^v, \\ dm_t &= \kappa_m(1 - m_t)dt + \sigma_m \sqrt{m_t} dW_t^m, \end{aligned} \quad (70)$$

⁶In the integral form, the covariance is

$$\int_0^t (\mu(v_s)^{\mathbb{M}} - \mu(v_s)^{\mathbb{Q}}) ds = \langle iu\sigma dW_{\mathcal{T}_t}, \sigma_v dW_{\mathcal{T}_t}^v \rangle = iu\sigma\sigma_v v_t \rho \mathcal{T}_t. \quad (65)$$

with W_t^m being an independent Brownian motion. The dynamics under measure $\mathbb{M}$ becomes,

$$\begin{aligned} dv_t &= (\kappa m_t - \kappa^{\mathbb{M}} v_t) dt + \sigma_v \sqrt{v_t} dW_t^v, \quad \kappa^{\mathbb{M}} = \kappa - iu\sigma\sigma_v\rho. \\ dm_t &= \kappa_m (1 - m_t) dt + \sigma_m \sqrt{m_t} dW_t^m. \end{aligned} \quad (71)$$

Writing the dynamics in a matrix notation with $V_t \equiv [v_t, m_t]^\top$, we have,

$$dV_t = \left(a - \kappa_V^{\mathbb{M}} V_t \right) dt + \sqrt{\Sigma V_t} dW_t^V, \quad (72)$$

with

$$a = \begin{bmatrix} 0 \\ \kappa_m \end{bmatrix}, \quad \kappa_V^{\mathbb{M}} = \begin{bmatrix} \kappa^{\mathbb{M}} & -\kappa \\ 0 & \kappa_m \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \sigma_v^2 & 0 \\ 0 & \sigma_m^2 \end{bmatrix}.$$

Given the two-factor affine structure for the activity rate dynamics, the Laplace transform in (61) is exponential affine in the current level of the two factors $V_0 = [v_0, m_0]^\top$:

$$\phi_s(u) = e^{iu(r-q)t - b(t)^\top V_0 - c(t)}, \quad (73)$$

where the coefficients $b(t)$ and $c(t)$ can be solved from a set of ordinary differential equations:

$$\begin{aligned} b'(t) &= \psi_x(u) b_V - (\kappa^{\mathbb{M}})^\top b(t) - \frac{1}{2} \Sigma [b(t) \odot b(t)], \\ c'(t) &= a^\top b(t), \end{aligned} \quad (74)$$

starting at $b(0) = 0$ and $c(0) = 0$, with $b_V = [1, 0]^\top$ denoting the instantaneous loading of the activity rate on the two factors and $\odot$ denoting the element-by-element product operation. The ordinary differential equations can be solved using standard numerical routines, such as an euler approximation or the fourth-order Runge-Kutta method.

When the return innovation is not driven by a diffusion, but by a pure jump Lévy process such as the one governed by the dampened power law in (12), we simply need to replace the characteristic exponent

of the concavity-adjusted diffusion component $\psi_x(u) = \frac{1}{2}(iu + u^2)\sigma^2$ by that of the concavity-adjusted jump component. With $\alpha \neq 0$ and $\alpha \neq 1$, we have:

$$\begin{aligned}\psi_x(u) = & -\Gamma(-\alpha)\lambda [(\beta_+ - iu)^\alpha - \beta_+^\alpha + (\beta_- + iu)^\alpha - \beta_-^\alpha] \\ & + iu\Gamma(-\alpha)\lambda [(\beta_+ - 1)^\alpha - \beta_+^\alpha + (\beta_- + 1)^\alpha - \beta_-^\alpha].\end{aligned}\tag{75}$$

We can also include both a diffusion and a jump component, in which case the characteristic exponent becomes the sum of the two. Most importantly, we can treat the specification of the Lévy process and the time change separately, and hence derive the characteristic exponent $\psi_x(u)$ and the Laplace transform separately. Therefore, we can combine any tractable Lévy specifications with any tractable activity rate dynamics, and the generalized Fourier transform for the resultant return dynamics is tractable.

5.2. Computing the Fourier inversions

With tractable solutions to the generalized Fourier transform of the return distribution, European option prices can be computed by inverting the Fourier transform. The literature considers two broad ways of inverting the transform. The first approach treat options analogous to a cumulative distribution function. Standard statistics books show how to invert the characteristic function to obtain a cumulative function. The inversion formula for option prices can be analogously proved. The second approach treats the option price analogous to a probability density function. In this case, for the transform to be well-defined, the characteristic exponent u in (57) needs to contain an imaginary component, the domains of which depend on the exact type of the payoff structure. Based on this analogy, option prices across the whole spectrum of strikes can be obtained via fast Fourier transform (FFT).

For both cases, the Fourier transforms for a wide variety of European payoffs can be obtained. Their values can then be obtained by inverting the corresponding transforms. Throughout the chapter, I use a European call option as an example to illustrate the transform methods. Indeed, in most situations, a call option value is all we need because most European payoff functions can be replicated by a portfolio of European call options across different strikes but at the same maturity.

The terminal payoff of the European call option at maturity t and strike K is,

$$\Pi_t = (S_t - K)1_{S_t \geq K}. \quad (76)$$

Since we have derived the Fourier transform of the asset returns, it is convenient to represent the payoff in log return terms,

$$\Pi_t = S_0(e^{\ln S_t/S_0} - e^{\ln K/S_0})1_{S_t \geq K} = S_0(e^{s_t} - e^k)1_{s_t \geq k}, \quad (77)$$

with $s_t = \ln S_t/S_0$ and $k = \ln K/S_0$. The time-0 value of the call option is then,

$$C(K, t) = S_0 e^{-rt} \mathbb{E}_0^{\mathbb{Q}} \left[(e^{s_t} - e^k) 1_{s_t \geq k} \right]. \quad (78)$$

Let $C(k) = C(K, t)/S_0$ denote the call option value in percentages of the current spot price level as a function of moneyness k and maturity t . In what follows, I focus on computing the relative call value $C(k)$. Then, we can simply multiply it by the spot price level to obtain the absolute call option value $C(K, t)$.⁷ We henceforth drop the maturity argument when no confusion shall occur.

5.2.1. The cumulative distribution analogy

We rewrite the call option value in terms of $x = -k$,

$$C(x) = C(k = -x) = e^{-rt} \mathbb{E}_0^{\mathbb{Q}} \left[(e^{s_t} - e^{-x}) 1_{-s_t \leq x} \right]. \quad (79)$$

Treating the call option value $C(x)$ analogous to a cumulative distribution, we define its transform as,

$$\chi_c(z) \equiv \int_{-\infty}^{\infty} e^{izk} dC(x), \quad z \in \mathbb{R}. \quad (80)$$

⁷Some broker dealers provide the relative percentage quote $C(k)$ instead of the absolute quote $C(K, t)$ to achieve quote stability by excluding the impact of spot price fluctuation.

We can derive this transform in terms of the Fourier transform of the return $\phi_s(u)$:

$$\begin{aligned}
\chi_c(z) &= e^{-rt} \mathbb{E}^{\mathbb{Q}} \left[\int_{-\infty}^{\infty} e^{izx} (e^s \delta_{-s \leq x} - e^{-x} \delta_{-s \leq x} + e^{-x} 1_{-s \leq x}) dx \right] \\
&= e^{-rt} \mathbb{E}^{\mathbb{Q}} \left[e^{(1-iz)s} - e^{(1-iz)s} + \int_{-s}^{\infty} e^{(iz-1)x} dx \right] \\
&= e^{-rt} \mathbb{E}^{\mathbb{Q}} \left[\frac{e^{(1-iz)s}}{(1-iz)} \right] = e^{-rt} \frac{\phi_s(-i-z)}{1-iz},
\end{aligned} \tag{81}$$

which is solved by first applying Fubini's theorem and then applying the result on the Fourier transform of a Dirac function $\delta_{-s \leq x}$. Thus, tractable forms for the return transform $\phi_s(u)$ also means tractable forms for the option transform $\chi_c(z)$.

Given this transform, the option value can be obtained via the following inversion formula:

$$C(x) = \frac{1}{2} \chi_c(0) + \frac{1}{2\pi} \int_0^{\infty} \frac{e^{izx} \chi_c(-z) - e^{-izx} \chi_c(z)}{iz} dz. \tag{82}$$

The inversion formula (and its proof) is very much analogous to the inversion formula for a cumulative distribution (Alan and Ord (1987)). The only difference is at the boundary: For a cumulative distribution, the transform evaluated at $z = 0$ is one; for the option transform, it is $\chi_c(0) = e^{-rt} \phi_s(-i) = e^{-qt}$. Given $C(x)$, we obtain $C(k) = C(k = -x)$. We can also directly define the inversion formula as,

$$C(k) = \frac{1}{2} \chi_c(0) + \frac{1}{2\pi} \int_0^{\infty} \frac{e^{-izk} \chi_c(-z) - e^{izk} \chi_c(z)}{iz} dz, \tag{83}$$

$$= e^{-rt} \left[\frac{1}{2} \phi_s(-i) - \frac{1}{2\pi} \int_0^{\infty} \left(e^{-izk} \frac{\phi_s(z-i)}{z^2 - iz} + e^{izk} \frac{\phi_s(-z-i)}{z^2 + iz} \right) dz \right]. \tag{84}$$

To compute the option value from the transform, the inversion formula in (82) asks for a numerical integration of an oscillating function. Fortunately, being a weighted average of cosines, the integrand exhibits much less oscillatory behavior than the transform $\psi(u)$ itself. The integral can numerically be evaluated using quadrature methods (Singleton (2001)).

Duffie, Pan, and Singleton (2000) and Leippold and Wu (2002) discuss the application of this approach for the valuation of general European-type state-contingent claims in the context of affine and quadratic models, respectively. In earlier works, e.g., Chen and Scott (1992), Heston (1993), Bates

(1996), and Bakshi, Cao, and Chen (1997), the call option value is often written as the portfolio of two contingent claims:

$$\begin{aligned} C(x) &= e^{-rt} \mathbb{E}^{\mathbb{Q}}[e^s] \frac{\mathbb{E}^{\mathbb{Q}}[e^s 1_{-s \leq x}]}{\mathbb{E}^{\mathbb{Q}}[e^s]} - e^{-rt} e^{-x} \mathbb{E}^{\mathbb{Q}}[1_{-s \leq x}] \\ &= e^{-qt} Q_1(x) - e^{-rt} e^{-x} Q_2(x), \end{aligned} \quad (85)$$

with $Q_1(x)$ and $Q_2(x)$ being the values of two contingent claims defined by,

$$Q_1(x) = \frac{\mathbb{E}^{\mathbb{Q}}[e^s 1_{-s \leq x}]}{\phi_s(-i)}, \quad Q_2(x) = \mathbb{E}^{\mathbb{Q}}[1_{-s \leq x}]. \quad (86)$$

Q_2 is simply the cumulative distribution of $-s$. Its transform is,

$$\chi_2(z) = \int_{-\infty}^{\infty} e^{izx} dQ_2(x) = \mathbb{E}^{\mathbb{Q}} \left[\int_{-\infty}^{\infty} e^{izx} \delta_{-s \leq x} dx \right] = \mathbb{E}^{\mathbb{Q}}[e^{-izs}] = \phi_s(-z). \quad (87)$$

The transform of $Q_1(x)$ is,

$$\chi_1(z) = \frac{1}{\phi_s(-i)} \mathbb{E}^{\mathbb{Q}} \left[\int_{-\infty}^{\infty} e^{izx} e^s \delta_{-s \leq x} dx \right] = \frac{1}{\phi_s(-i)} \mathbb{E}^{\mathbb{Q}}[e^{(1-iz)s}] = \frac{\phi_s(-z-i)}{\phi_s(-i)}. \quad (88)$$

Applying the inversion formula in (83), we have the values for the two contingent claims as,

$$Q_1(k) = \frac{1}{2} + \frac{1}{2\pi\phi_s(-i)} \int_0^{\infty} \frac{e^{-izk} \phi_s(z-i) - e^{izk} \phi(-z-i)}{iz} dz, \quad (89)$$

$$Q_2(k) = \frac{1}{2} + \frac{1}{2\pi} \int_0^{\infty} \frac{e^{-izk} \phi_s(z) - e^{izk} \phi_s(-z)}{iz} dz. \quad (90)$$

Doing one numerical integration according to (84) should be more efficient than doing two numerical integrations according to (89) and (90).

5.2.2. The probability density analogy

The second approach treats the option price analogous to a probability density and defines the Fourier transform of the option value as,

$$\chi_p(z) \equiv \int_{-\infty}^{\infty} e^{izk} C(k) dk, \quad z = z_r - iz_i, \quad z_r \in \mathbb{R}, z_i \in \mathcal{D} \subseteq \mathbb{R}^+. \quad (91)$$

The transform coefficient z is extended to the complex plane is to guarantee the finiteness of the transform. For the call option value, the transform is,

$$\begin{aligned} \chi_p(z) &= \int_{-\infty}^{\infty} e^{izk} \mathbb{E}^{\mathbb{Q}} \left[e^{-rt} (e^s - e^k) 1_{s \geq k} \right] dk = e^{-rt} \mathbb{E}^{\mathbb{Q}} \left[\int_{-\infty}^{\infty} e^{izk} (e^s - e^k) 1_{s \geq k} dk \right] \\ &= e^{-rt} \mathbb{E}^{\mathbb{Q}} \left[\int_{-\infty}^s e^{izk} (e^s - e^k) dk \right] = e^{-rt} \mathbb{E}^{\mathbb{Q}} \left[\left(\frac{e^{izk} e^s}{iz} - \frac{e^{(iz+1)k}}{iz+1} \right) \Big|_{k=-\infty}^{k=s} \right]. \end{aligned} \quad (92)$$

For $e^{izk} = e^{iz_r k + z_i k}$ to be convergent (to zero) at $k = -\infty$, we need $z_i > 0$, under which $e^{(iz+1)k}$ also converges to zero.⁸ With $z_i > 0$, the transform for the call option value becomes,

$$\chi_p(z) = e^{-rt} \mathbb{E}^{\mathbb{Q}} \left[\frac{e^{(1+iz)s}}{iz} - \frac{e^{(iz+1)s}}{iz+1} \right] = e^{-rt} \frac{\phi_s(z-i)}{(iz)(iz+1)}. \quad (93)$$

For some return distributions, the return transform $\phi_s(z-i) = \mathbb{E}^{\mathbb{Q}}[e^{(1+iz)s}]$ is well-defined only when z_i is in a subset of the real line. In equation (91), we use $\mathcal{D} \in \mathbb{R}^+$ to denote the subset that both guarantees the convergence of e^{izk} and $e^{(iz+1)k}$ at $k = -\infty$, and assures the finiteness of the transform $\phi_s(z-i)$.

Given a finite transform $\chi_p(z)$ for the call option, the option value can be computed from the following inversion formula:

$$C(k) = \frac{1}{2} \int_{-iz_i - \infty}^{-iz_i + \infty} e^{-izk} \chi_p(z) dz = \frac{e^{-z_i k}}{\pi} \int_0^{\infty} e^{-iz_r k} \chi_p(z_r - iz_i) dz_r. \quad (94)$$

⁸For other types of contingent claims, the transform will take different forms and the required domain for z_i that guarantees the finiteness of the transform varies accordingly.

We can approximate the integral using summations:

$$C(k) \approx \widehat{C}(k) = \frac{e^{-z_i k}}{\pi} \sum_{n=0}^{N-1} e^{-iz_r(n)k} \chi_p(z_r(n) - iz_i) \Delta z_r, \quad (95)$$

where $z_r(n)$ are the nodes of z_r and Δz_r is the spacing between nodes. The fast Fourier transform (FFT) is an efficient algorithm for computing the discrete Fourier coefficients. The discrete Fourier transform is a mapping of $\mathbf{f} = (f_0, \dots, f_{N-1})^\top$ on the vector of Fourier coefficients $\mathbf{d} = (d_0, \dots, d_{N-1})^\top$, such that

$$d_j = \sum_{n=0}^{N-1} f_n e^{-jn \frac{2\pi}{N} i}, \quad j = 0, 1, \dots, N-1. \quad (96)$$

We use $\mathbf{d} = D(\mathbf{f})$ to denote the fast Fourier transform, which allows the efficient calculation of $\mathbf{d}$ if N is an even number, say $N = 2^m$, $m \in \mathbb{N}$. The algorithm reduces the number of multiplications in the required N summations from an order of 2^{2m} to that of $m2^{m-1}$, a very considerable reduction. By a suitable choice of Δz_r and a discretization scheme for k , we can cast the approximation in the form of (96) to take advantage of the computational efficiency of the FFT.

Following Carr and Madan (1999), we set $z_r(n) = \eta n$ and $k_j = -b + \lambda j$, and require $\eta\lambda = 2\pi/N$. Then, we can cast the option valuation approximation in (95) in the form of the FFT summation in (96):

$$\widehat{C}(k_j) = \sum_{n=0}^{N-1} f_n e^{-jn \frac{2\pi}{N} i} = D_j(\mathbf{f}), \quad j = 0, 1, \dots, N-1, \quad (97)$$

with

$$f_n = \frac{1}{\pi} e^{-z_i k_j + i b \eta n} \eta \chi_p(\eta n - iz_i). \quad (98)$$

Under such a discretization scheme, the effective upper limit for the integration is $N\eta$, with a spacing of η . The range of log strike level is from $-b$ to $N\lambda - b$, with a uniform spacing of λ in the log strike. To put at-the-money ($k = 0$) option at the middle of the strike range, we can set $b = N\lambda/2$.

The restriction of $\eta\lambda = 2\pi/N$ reveals the trade-off between a fine grid in log strike and a fine grid in summation. With $N = 2^{12}$, Carr and Madan (1999) set $\eta = 0.25$ to price stock options. To price currency and interest-rate options, I often set $\eta = 1$ to generate a finer spacing of strikes and hence more option values within the relevant range. The choice of the imaginary part of the transform coefficient z_i also affects the numerical accuracy of the fast Fourier inversion. Lee (2004) provides detailed analysis

on the error bounds and on the choice of the imaginary part of the transform coefficient z_i . Carr and Madan (1999) suggest to incorporate Simpson's rule to improve the integration accuracy with a fixed spacing η , but I find that in many applications incorporating Simpson's rule does not improve the integration accuracy.

5.2.3. Fractional FFT

Recently, Chourdakis (2005) adopts the fractional Fourier transform (FRFT) method of Bailey and Swartztrauber (1991) in inverting the option transform $\chi_p(z)$. The method can efficiently compute,

$$d_j = \sum_{n=0}^{N-1} f_n e^{-jn\alpha i}, \quad j = 0, 1, \dots, N-1, \quad (99)$$

for any value of the parameter α . The standard FFT can be seen as a special case for $\alpha = 2\pi/N$. Therefore, we can use the FRFT method to compute,

$$\widehat{C}(k, t) = \sum_{n=0}^{N-1} f_n e^{-jn\eta\lambda i}, \quad j = 0, 1, \dots, N-1, \quad (100)$$

with without the trade-off between the summation grid η and the strike spacing λ .

We use $\mathbf{d} = D(\mathbf{f}, \alpha)$ to denote the FRFT operation, with $D(\mathbf{f}) = D(\mathbf{f}, 2\pi/N)$ being the standard FFT as a special case. An N -point FRFT can be implemented by invoking three $2N$ -point FFT procedures. Define the following $2N$ -point vectors:

$$\mathbf{y} = \left(\left(f_n e^{i\pi n^2 \alpha} \right)_{n=0}^{N-1}, (0)_{n=0}^{N-1} \right), \quad (101)$$

$$\mathbf{z} = \left(\left(e^{i\pi n^2 \alpha} \right)_{n=0}^{N-1}, \left(e^{i\pi (N-n)^2 \alpha} \right)_{n=0}^{N-1} \right). \quad (102)$$

The FRFT is given by,

$$D_k(\mathbf{h}, \alpha) = \left(e^{i\pi k^2 \alpha} \right)_{k=0}^{N-1} \odot D_k^{-1} (D_j(\mathbf{y}) \odot D_j(\mathbf{z})), \quad (103)$$

where $D_k^{-1}(\cdot)$ denotes the inverse FFT operation and $\odot$ denotes element-by-element vector multiplication. Due to the multiple application of the FFT operations, Chourdakis (2005) shows that an N -point FRFT procedure demands a similar number of elementary operations as an $4N$ -point FFT procedure. However, given the free choices on λ and η , FRFT can be applied more efficiently. Using a smaller N with FRFT can achieve the same option pricing accuracy as using a much larger N with FFT. Numerical analysis shows that with the similar computational time, the FRFT method can often achieve better computational accuracy than the FFT method. The accuracy improvement is larger when we have a better understanding of the model and model parameters so that we can set the boundaries more tightly. Nevertheless, the analysis also reveals a few cases of complete breakdown when the model takes extreme parameters and when the bounds are set too tight. Hence, the more freedom also asks for more discretion and caution in applying this method to generate robust results in all situations. This concern becomes especially important for model estimation, during which the trial model parameters can vary greatly.

6. Estimating Lévy processes with and without time changes

Estimation can be classified into three categories: (1) Estimating a statistical process to capture the behavior of the time-series returns, (2) estimating a risk-neutral process to match the option price behavior, and (3) jointly estimating the statistical and risk-neutral process using both time-series returns and option prices and learning the behavior of market prices on different sources of risks.

6.1. Estimating statistical dynamic using time-series returns

Without time change, a Lévy process implies that the security returns are iid. Thus, we can regard each day's return as random draws from the same distribution. This property makes the maximum likelihood method easy to implement. For the Lévy processes that I have discussed in this paper, only a few of them have analytical density functions, but virtually all of them have analytical characteristic functions. We can use fast Fourier transform (FFT) to numerically convert the characteristic function into density functions. Carr, Geman, Madan, and Yor (2002) use this method to estimate the CGMY

model to stock returns. To implement this method, we normally need to use a large number N for the FFT so that we obtain numerical density values at a fine grid of realizations. Then, we can map the actual data to the grids by grouping the actual realizations into different bins that match the grids of the FFT and assign the same likelihood for realizations within the same bin. Alternatively, we can simply interpolate the density values from the FFT to match the actual realizations. Furthermore, to improve numerical stability and to generate enough points in the relevant FFT range, it is often helpful to standardize the return series (Wu (2006)).

The estimation becomes more involved when the model contains random time changes. Since the activity rates are not observable, some filtering technique is often necessary to determine the current level of the activity rates. Eraker, Johannes, and Polson (2003) and Li, Wells, and Yu (2004) propose to estimate the dynamics using a Bayesian approach involving Markov Chain Monte Carlo (MCMC) simulation. They use MCMC to Bayesian update the distribution of both the state variables and model parameters. Javaheri (2005) propose a maximum likelihood method in estimating time-changed Lévy processes. Under this method, the distribution of the activity rates are predicted and updated according to Bayesian rules and using Markov Chain Monte Carlo simulation. Then, the model parameters are estimated by maximizing the likelihood of the time-series returns. Kretschmer and Pigorsch (2004) propose to use the efficient method of moments (EMM) of Gallant and Tauchen (1996).

6.2. *Estimating risk-neutral dynamic to fit option prices*

If the objective is to estimate a Lévy return risk-neutral process to option prices, nonlinear least square or some variant of it is the most direct method to use. Since a Lévy process implies iid returns, the conditional return distribution over a fixed time horizon remains the same at different dates. Accordingly, the option price behavior across strikes and time-to-maturities, when scaled by the spot price, should remain the same across the different dates. In particular, the Black-Scholes implied volatility surface across moneyness and time-to-maturity should remain the same across different days. In reality, however, the option price behavior does change over time. For example, the implied volatility levels vary over time. The shape of the implied volatility smile also varies over time. A Lévy model without time change cannot capture these time variations. A common practice in the industry is to estimate the model daily, that is, to use different model parameters to match the different implied volatility levels

and shapes at different days. This method is convenient and is also used in early academic works, e.g., Bakshi, Cao, and Chen (1997) and Carr and Wu (2003a).

In fact, even for one day, most Lévy processes have difficulties fitting the implied volatility surface across different maturities. The implied volatility smile observed from the market often persists as maturity increases, implying that the risk-neutral return distribution remains highly non-normal at long horizons. Yet, since Lévy models imply iid returns, if the return variance is finite under the model specification, the classic central limit theorem dictates that the skewness of the return distribution declines like the reciprocal of the square root of the horizon and the excess kurtosis declines like the reciprocal of horizon. Hence, return non-normality declines rapidly with increasing maturities. For these models, calibration is often forced to be done at each maturity. A different set of model parameters are used to fit the implied volatility smile at different maturities.

Carr and Wu (2003a) uses a maximum negatively skewed α -stable process to model the stock index return. Although the model-implied return distribution is iid, the model-implied return variance is infinite and hence the central limit theorem does not apply. Thus, the model is capable of generating persistent implied volatility smiles across maturities. Wu (2006) use the dampened power law to model the index return innovation. With exponential dampening under the statistical measure, return variance is finite and the central limit theorem applies. The statistical return distribution is non-normal at high sampling frequencies but converges to normal rapidly with time aggregation. However, by applying a measure change using an exponential martingale, the dampening on the left tail can be made to disappear under the risk-neutral measure so that the return variance becomes infinite under the risk-neutral measure and the return non-normality no longer disappears with increasing option maturity.

Applying stochastic time change to Lévy processes not only generates time variation in the return distribution, but also generates cross-sectional option price behaviors that are more consistent with market observations. For example, a persistent activity rate process can generate non-normality out of a normal return innovation and can slow down the convergence of a non-normal return distribution to normality. For daily calibration, the unobservable activity rates are treated the same as model parameters. They are all used as free inputs to make the model values fit market observations.

A dynamically consistent estimation is to keep the model parameters constant and only allow the activity rates to vary over time. Huang and Wu (2004) employ a nested nonlinear least square procedure for this purpose. Given parameter guesses, they minimize the pricing errors at each day to infer the activity rates at that day. Then, the parameters are chosen to minimize the aggregate pricing errors over the whole sample period. Carr and Wu (2004a) cast the models in a state-space form and estimate the model parameters using the maximum likelihood method. The state propagation equations are defined by the time-series dynamics of the activity rates and the measurement equations are defined on the option prices. Given parameter guesses, they use an extended version of the Kalman filter, the unscented Kalman filter (Wan and van der Merwe (2001)), to obtain the forecasts and filtering on the conditional mean and variance of the states and measurements. Then, they construct the likelihood of the option series assuming normally distributed forecasting errors.

6.3. Static and dynamic consistency in model estimation

Daily calibration or calibration at each option maturity raises the issue of internal consistency. Option values generated from a no-arbitrage model are international consistent with one another and do not generate arbitrage opportunities among themselves. When a model is re-calibrated at each maturity, the option values generated at different maturities are essentially from different models and hence the international consistency between them is no longer guaranteed. When a model is calibrated daily, option values generated from the model at one day are not guaranteed to be consistent with option values generated at another day. One of the potential dangers of doing daily calibration is in risk management. A “fully” hedged option portfolio based on a model assuming constant model parameters is destined to generate hedging errors if the model parameters are altered on a daily basis.

Both the academia and practitioners appreciate the virtue of being both cross-sectionally and dynamically consistent. Nevertheless, building a dynamically consistent model that fits the market data well can be difficult. Hence, the daily calibration method can be regarded as a compromise. It remains true that a hedging strategy with constant parameter assumptions is bound to generate hedging errors when the model parameters are altered. One way to minimize the impact of varying parameters is to consider very short investment horizons in the hope that the hedging errors due to dynamic inconsistency are small over a short time interval. For example, market makers can be regarded as very

short-term investors as they rarely take long-term inventories and hence rarely take long-term views in providing two-side quotes to the market. Therefore, dynamic consistency may not be as an overriding concern as it is to long-term investors. The more pressing concern for market makers is to achieve cross-sectional consistency across quotes for different contracts at a point in time. Furthermore, since they need to provide two-sided quotes, they often need a model that can match the current market quotes well.

On the other hand, for a hedge fund that bets on long-term convergence, a model that always fits the data well is not the key requirement. In fact, since their objective is to find market mispricings, it is important that their model can generate values that differ from the market sometimes. A good model produces pricing errors that are zero on average and transient in nature, so that if the model picks out a security that is over-valued, the over-valuation disappears in the near future. However, although they have a less stringent requirement on the model's fitting performance, they often have a more stringent requirement for dynamic consistency when they bet on long-term convergence. To them, it is important to keep the model parameters fixed over time and only allow state variables to vary, even if such a practice increases the model complexity and sometimes also increase the pricing errors of the model.

In a dynamically consistent model, the parameters that are allowed to vary daily should be converted into state variables, and their dynamics should be priced when valuing a contingent claim. Random time changes provide an intuitive and tractable way of turning a static model to a dynamic one. It is no longer an ideal that we can build tractable models that can generate reasonable pricing performance while maintaining dynamic consistency. Recent developments in econometrics further enables us to estimate these models with dynamic consistency constraints and within a reasonable time framework. Once estimated, updating the activity rates based on newly arrived option quotes can be done almost instantaneously. Hence, it causes no delays in trading or market making.

6.4. Joint estimation of statistical and risk-neutral dynamics

One of the frontiers in the academic literature is to exploit the information in the derivatives market to infer the market prices on various sources of risks. While a long time series can be used to estimate the statistical dynamics of the security return, a large cross-section of option prices across multiple

strikes and maturities provide important information about the risk-neutral dynamics. The market prices of various sources of risks dictate the difference between the return dynamics under the two measures. Hence, estimation using both time series and cross-sectional data can help us identify the dynamics under both measures and the market pricing on various sources of risks.

Pan (2002) uses the generalized methods of moments to estimate a jump-diffusion stochastic volatility model under both probability measures and study the jump risk premia implicit in options. The moment conditions are constructed using both options and time-series returns. Eraker (2004) estimate similar dynamics under both measures using the MCMC approach. At each day, he uses the time-series returns and a few randomly sampled option prices. As a result, many available options data are thrown out in his estimation. Bakshi and Wu (2005) propose a maximum likelihood approach, where the likelihood on options and on time-series returns are constructed sequentially and the maximization is over the sum of the likelihoods on the two sets of data. First, they cast the activity rate dynamics into a state-propagation equation and the option prices into measurement equations. Second, they use an extended version of a Kalman filter to predict and update on the activity rates. Third, the likelihood on the options are constructed based on the forecasting errors on the options assuming normal forecasting errors. Fourth, they take the filtered activity rates as given and construct the likelihood of the returns conditional on the filtered activity rates. The conditional likelihood can be obtained using fast Fourier inversion of the conditional characteristic function. Finally, model parameters are chosen to maximize the sum of the likelihood of the time-series returns and option prices. They use this estimation procedure to analyze the variation of various sources of market prices around the Nasdaq bubble period.

7. Concluding remarks

Lévy processes with and without time changes have become the universal building blocks for financial security returns. Different Lévy components can be used to capture both continuous and discontinuous movements. Stochastic time changes can be applied to randomize the intensity of these different movements to generate stochastic time variation in volatility and even in higher return moments. In this chapter, I provide a summary on how different return behaviors can be captured by different Lévy

components and different ways of applying time changes, under both the risk-neutral measure and the statistical measure. I also discuss how to compute European option values under these specifications using Fourier transform methods, and how to estimate the model parameters using time-series returns and/or option prices.

References

- Aït-Sahalia, Y., 2004. Disentangling diffusion from jumps. *Journal of Financial Economics* 74, 487–528.
- Alan, S., Ord, J. K., 1987. *Kendall's Advanced Theory of Statistics Vol. 1*. Oxford University Press, New York.
- Ané, T., Geman, H., 2000. Order flow, transaction clock and normality of asset returns. *Journal of Finance* 55, 2259–2284.
- Backus, D., Foresi, S., Mozumdar, A., Wu, L., 2001. Predictable changes in yields and forward rates. *Journal of Financial Economics* 59, 281–311.
- Backus, D., Foresi, S., Telmer, C., 2001. Affine term structure models and the forward premium anomaly. *Journal of Finance* 56, 279–304.
- Bailey, D. H., Swartztrauber, P. N., 1991. The fractional Fourier transform and applications. *SIAM Review* 33, 389–404.
- Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52, 2003–2049.
- Bakshi, G., Ju, N., Ou-Yang, H., 2004. Estimation of continuous-time models with an application to equity volatility. Working paper. University of Maryland.
- Bakshi, G., Wu, L., 2005. Investor irrationality and the nasdaq bubble. Working paper. University of Maryland and Baruch College.
- Balduzzi, P., Das, S., Foresi, S., 1998. The central tendency: A second factor in bond yields. *Review of Economics and Statistics* 80, 62–72.
- Barndorff-Nielsen, O. E., 1998. Processes of normal inverse Gaussian type. *Finance and Stochastics* 2, 41–68.
- Bates, D., 1996. Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies* 9, 69–107.
- Bertoin, J., 1996. *Lévy Processes*. Cambridge University Press, Cambridge.

- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–654.
- Broadie, M., Chernov, M., Johannes, M., 2002. Jumps in volatility: The evidence of index options. Working paper. Columbia University.
- Carr, P., Geman, H., Madan, D., Yor, M., 2002. The fine structure of asset returns: An empirical investigation. *Journal of Business* 75, 305–332.
- Carr, P., Madan, D., 1999. Option valuation using the fast Fourier transform. *Journal of Computational Finance* 2, 61–73.
- Carr, P., Sun, J., 2005. A new approach for option pricing under stochastic volatility. Working paper. Bloomberg LP.
- Carr, P., Wu, L., 2003a. Finite moment log stable process and option pricing. *Journal of Finance* 58, 753–777.
- Carr, P., Wu, L., 2003b. What type of process underlies options? A simple robust test. *Journal of Finance* 58, 2581–2610.
- Carr, P., Wu, L., 2004a. Stochastic skew in currency options. Working paper. New York University and Baruch College.
- Carr, P., Wu, L., 2004b. Time-changed Lévy processes and option pricing. *Journal of Financial Economics* 71, 113–141.
- Carr, P., Wu, L., 2005a. Fiat money as stock in small open economies: Theory and evidence on sovereign credit default swap and currency options. Working paper. New York University and Baruch College.
- Carr, P., Wu, L., 2005b. Stock options and credit default swaps: A joint framework for valuation and estimation. Working paper. New York University and Baruch College.
- Chacko, G., Viceira, L., 2003. Spectral GMM estimation of continuous-time processes. *Journal of Econometrics* 116, 259–292.
- Chen, R.-R., Scott, L., 1992. Pricing interest rate options in a two-factor cox-ingersoll-ross model of the term structure. *Review of Derivatives Studies* 5, 613–636.

- Cheridito, P., Filipović, D., Kimmel, R. L., 2003. Market price of risk specification for affine models: Theory and evidence. Working paper. Princeton University.
- Chourdakis, K. M., 2005. Option pricing using the fractional FFT. *Journal of Computational Finance* 8, 1–18.
- Cox, J. C., Ingersoll, J. E., Ross, S. R., 1985. A theory of the term structure of interest rates. *Econometrica* 53, 385–408.
- Dai, Q., Singleton, K., 2002. Expectation puzzles, time-varying risk premia, and affine models of the term structure. *Journal of Financial Economics* 63, 415–441.
- David, A., Veronesi, P., 1999. Option prices with uncertain fundamentals: Theory and evidence on the dynamics of implied volatilities. Working paper. Federal Reserve and University of Chicago.
- Duffee, G. R., 2002. Term premia and interest rate forecasts in affine models. *Journal of Finance* 57, 405–443.
- Duffie, D., Filipović, D., Schachermayer, W., 2003. Affine processes and applications in finance. *Annals of Applied Probability* 13, 984–1053.
- Duffie, D., Kan, R., 1996. A yield-factor model of interest rates. *Mathematical Finance* 6, 379–406.
- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump diffusions. *Econometrica* 68, 1343–1376.
- Duffie, D., Pedersen, L. H., Singleton, K., 2003. Modeling sovereign yield spreads: A case study of Russian debt. *Journal of Finance* 58, 119–160.
- Duffie, D., Singleton, K., 1999. Modeling term structure of defaultable bonds. *Review of Financial Studies* 12, 687–720.
- Duffie, D., Singleton, K., 2003. *Credit Risk: Pricing, Measurement and Management*. Princeton University Press, Princeton, NJ.
- Eberlein, E., Keller, U., Prause, K., 1998. New insights into smile, mispricing, and value at risk: The hyperbolic model. *Journal of Business* 71, 371–406.
- Engle, R., 2004. Risk and volatility: Econometric models and financial practice. *American Economic Review* 94, 405–420.

- Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59, 1367–1404.
- Eraker, B., Johannes, M., Polson, N., 2003. The impact of jumps in equity index volatility and returns. *Journal of Finance* 58, 1269–1300.
- Fama, E. F., 1965. The behavior of stock market prices. *Journal of Business* 38, 34–105.
- Foresi, S., Wu, L., 2005. Crash-o-phobia: A domestic fear or a worldwide concern?. *Journal of Derivatives* forthcoming.
- Gallant, A. R., Tauchen, G., 1996. Which moment to match. *Econometric Theory* 12, 657–681.
- Heston, S., 1993. Closed-form solution for options with stochastic volatility, with application to bond and currency options. *Review of Financial Studies* 6, 327–343.
- Heston, S., 1997. A simple new formula for options with stochastic volatility. Working paper. University of Maryland.
- Huang, J., Wu, L., 2004. Specification analysis of option pricing models based on time-changed Lévy processes. *Journal of Finance* 59, 1405–1440.
- Ishida, I., Engle, R. F., 2002. Modeling variance of variance: The square root, the affine, and the CEV GARCH models. Working paper. New York University.
- Jacod, J., Shiryaev, A. N., 1987. *Limit Theorems for Stochastic Processes*. Springer-Verlag, Berlin.
- Javaheri, A., 2005. *Inside Volatility Arbitrage; The Secrets of Skewness*. John Wiley & Sons, London.
- Johnson, T. C., 2002. Volatility, momentum, and time-varying skewness in foreign exchange returns. *Journal of Business and Economic Statistics* 20, 390–411.
- Jones, C. S., 2003. The dynamics of stochastic volatility: Evidence from underlying and options markets. *Journal of Econometrics* 116, 181–224.
- Kou, S. G., 2002. A jump-diffusion model for option pricing. *Management Science* 48, 1086–1101.
- Kou, S. G., Wang, H., 2004. Option pricing under a double-exponential jump diffusion model. *Management Science* 50, 1178–1192.
- Kretschmer, U., Pigorsch, C., 2004. EMM estimation of time-changed Lévy processes. Working paper. University of Bonn and University of Munich.

- Küchler, U., Sørensen, M., 1997. Exponential Families of Stochastic Processes. Springer, New York.
- Lando, D., 1998. On Cox processes and credit risky securities. *Review of Derivatives Research* 2, 99–120.
- Lee, R. W., 2004. Option pricing by transform methods: Extensions, unification and error control. *Journal of Computational Finance* 7, 51–86.
- Leippold, M., Wu, L., 2002. Asset pricing under the quadratic class. *Journal of Financial and Quantitative Analysis* 37, 271–295.
- Leippold, M., Wu, L., 2003. Multi-currency quadratic models: Theory and evidence. manuscript. University of Zurich and Fordham University.
- Lewis, A. L., 2000. Option Valuation under Stochastic Volatility. Finance Press, Newport Beach, California, USA.
- Lewis, A. L., 2001. A simple option formula for general jump-diffusion and other exponential Lévy processes. Working paper. Envision Financial Systems and OptionCity.net Newport Beach, California, USA.
- Li, H., Wells, M., Yu, L., 2004. A MCMC analysis of time-changed lévy processes of stock return dynamics. Working paper. Cornell University.
- Madan, D., Seneta, E., 1990. The variance gamma (V.G.) model for share market returns. *Journal of Business* 63, 511–524.
- Madan, D. B., Carr, P. P., Chang, E. C., 1998. The variance gamma process and option pricing. *European Finance Review* 2, 79–105.
- Mandelbrot, B. B., 1963. The variation of certain speculative prices. *Journal of Business* 36, 394–419.
- Medvedev, Scaillet, O., 2003. A simple calibration procedure of stochastic volatility models with jumps by short term asymptotics. Working paper. HEC, University of Geneva.
- Merton, R. C., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.
- Mo, H., Wu, L., 2005. International capital asset pricing: Theory and evidence from index options. Working paper. Baruch College.

- Monroe, I., 1978. Processes that can be embedded in Brownian motion. *Annals of Probability* 6, 42–56.
- Pan, J., 2002. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63, 3–50.
- Pan, J., Singleton, K. J., 2005. Default and recovery implicit in the term structure of sovereign CDS spreads. Working paper. Stanford University and MIT.
- Roberds, W., Whiteman, C., 1999. Endogenous term premia and anomalies in the term structure of interest rates: Explaining the predictability smile. *Journal of Monetary Economics* 44, 555–580.
- Santa-Clara, P., Yan, S., 2005. Crashes, volatility, and the equity premium: Lessons from S&P 500 options. Working paper. UCLA.
- Sato, K.-I., 1999. *Lévy Processes and Infinitely Divisible Distributions*. Cambridge University Press, Cambridge.
- Schoutens, W., 2003. *Lévy Processes in Finance: Pricing Financial Derivatives*. John Wiley & Sons Ltd, London.
- Singleton, K. J., 2001. Estimation of affine asset pricing models using the empirical characteristic function. *Journal of Econometrics* 102, 111–141.
- Titchmarsh, E. C., 1986. *Introduction to the Theory of Fourier Integrals*. Chelsea Publication Co., New York.
- Wan, E. A., van der Merwe, R., 2001. The unscented Kalman filter. In: Haykin, S. (Eds.), *Kalman Filtering and Neural Networks*. Wiley & Sons Publishing, New York.
- Wu, L., 2005. Variance dynamics: Joint evidence from options and high-frequency returns. Working paper. Baruch College.
- Wu, L., 2006. Dampened power law: Reconciling the tail behavior of financial asset returns. *Journal of Business* 79, forthcoming.