Dampened Power Law:
Reconciling the Tail Behavior of Financial Security Returns*

Liuren Wu[†]

Zicklin School of Business, Baruch College

First draft: May 31, 2001

This version: February 13, 2004

*I thank Turan Bali, Peter Carr, Xiong Chen, Jingzhi Huang, Albert Madansky (the editor), and an anonymous referee for comments. All remaining errors are mine.

[†]One Bernard Baruch Way, Box B10-254, New York, NY 10010; tel: (646) 312-3509; fax: (646) 312-3451; liuren_wu@baruch.cuny.edu; <http://faculty.baruch.cuny.edu/lwu/>.

Dampened Power Law:

Reconciling the Tail Behavior of Financial Security Returns

ABSTRACT

This paper proposes a stylized model that reconciles several seemingly conflicting findings on financial security returns and option prices. The model is based on a pure jump Lévy process, wherein the jump arrival rate obeys a power law dampened by an exponential function. The model allows for different degrees of dampening for positive and negative jumps, and also different pricing for upside and downside market risks. Calibration of the model to the S&P 500 index shows that the market charges only a moderate premium on upward index movements, but the maximally allowable premium on downward index movements.

JEL CLASSIFICATION CODES: E43, G12, G13, C51.

KEY WORDS: dampened power law; α -stable distribution; central limit theorem; upside movement; downside movement.

Dampened Power Law:

Reconciling the Tail Behavior of Financial Security Returns

I. Introduction

In the early 1960s, Mandelbrot (1963) and Fama (1965) found that stock return distributions possess power tails that are invariant to time aggregation and scaling. Such findings led them to believe that these returns should follow an α -stable distribution, rather than the commonly assumed Gaussian distribution.

As a test of the stable law, several studies investigate the stability-under-addition property of asset returns. These studies find that in most cases, asset returns do converge to normality with time aggregation, contradicting the implication of an α -stable distribution. Examples of such studies include Barnea and Downes (1973), Brenner (1974), Fielitz and Rozelle (1983), Haggeman (1978), Hall, Brorsen, and Irwin (1989), Hsu, Miller, and Wichern (1974), Officier (1972), and Teichmoeller (1971).

Carr and Wu (2003) use options on the S&P 500 index to investigate how the risk-neutral return distribution for the equity index varies with the time horizon. They find that the risk-neutral distribution for the equity index return is highly non-normal, and this return non-normality does not decline with increasing time horizon, supporting the stability-under-addition property of an α -stable distribution.

These different pieces of evidence are all robust findings about the financial market, but seemingly contradict one another, adding fuel to the decade-long debate on whether an α -stable distribution is a realistic modeling choice for asset returns. An α -stable distribution captures the power law decay of

the tails of the return distribution, generates the risk-neutral stability-under-addition property observed from the options data, but is inconsistent with the time series evidence that asset returns converge to normality with time aggregation under the objective measure.

In this paper, I propose a stylized model that reconciles the seemingly conflicting pieces of evidence. The model generates power tails for asset returns to match the evidence on the power law. To guarantee that the central limit theorem holds under the objective measure, I dampen the power tails by an exponential function. The dampening is sufficient to guarantee finite return moments and the applicability of the central limit theorem, but not enough to overrule the power decay of the tails. I label this model as the exponentially **dampened power law** (DPL).

To link the time series behavior of the asset return to its risk-neutral behavior inferred from the options data, I propose a measure change defined by an extended exponential martingale. The measure change allows the market to price downside and upside risks differently. As a special example, when the market charges the maximally allowable premium by no arbitrage on downside risk, the dampening on the left tail of the return distribution disappears under the risk-neutral measure. As the central limit theorem no longer applies without the dampening on the left tail, the risk-neutral return distribution shows stability under addition, even though the objective return distribution does not.

I calibrate the model to the S&P 500 index returns and the index option prices. The calibration exercise sheds light on the market's distinct treatment of downside and upside index movements. The calibration results show that although the market participants charge only a moderate premium on upside movements in the equity index, they charge the maximally allowable premium on downside index movements.

The DPL specification applies not only to the equity market, but also to currencies. The time series behaviors of equity and currency returns are similar, but their respective options markets exhibit quite distinctive behaviors. First, the option implied risk-neutral distribution for the equity index returns is highly skewed to the left, but the risk-neutral distribution for currency returns is relatively symmetric. Second, the non-normality of the equity index return risk-neutral distribution does not decline as option maturity increases, but the non-normality of the currency return inferred from the currency options market declines steadily as predicted by the central limit theorem. Under the framework of the dampened power law, these differences imply that the market participants do not distinguish the direction of the currency movement, although they distinguish the direction of the equity index movement.

The DPL specification reconciles a series of seemingly conflicting evidence concerning financial security returns. Nevertheless, I do not regard the DPL as the final answer for modeling, but rather as a springboard for more comprehensive modeling endeavors. As an illustration, I show how the pure jump component underlying the DPL model can be tightly knitted with an additional diffusion component and stochastic volatility.

The most germane to my work is the CGMY model of Carr, Geman, Madan, and Yor (2002). Although they derive the CGMY model by extending the variance-gamma specification in Madan and Seneta (1990) and Madan, Carr, and Chang (1998), the CGMY model follows the exponentially dampened power law. Carr, Geman, Madan, and Yor consider the application of the CGMY model both in modeling the time series property of equity returns and in pricing equity and equity index options. Carr, Geman, Madan, and Yor (2003) extends the model to incorporate stochastic volatility. Compared to their work, the key contribution of my work in this paper lies in the documentation and reconciliation of the major stylized evidence defining the tail behavior of financial security returns. The DPL specifi-

cation is also related to the physics literature on truncated Lévy flights (Mantegna and Stanley (1995)). Boyarchenko and Levendorskii (2000) consider option pricing under such processes.

Another contribution of this paper is my distinct treatment of upside and downside market movements for pricing. The sharp difference in the equity index return distribution under the objective and the risk-neutral measures has attracted great attention and curiosity from the academia. Jackwerth (2000) and Engle and Rosenberg (2002) find that, to reconcile the index return distribution under the two measures, one may end up with some oddly shaped preferences, with sections that are locally risk-loving rather than risk averse. Bates (2001) tries to explain the difference in an equilibrium model. My distinct treatment of downside and upside risks not only reconciles the difference in the asymmetry of the index return distribution under the two measures, but also explains their different behaviors along the maturity dimension. By charging the maximally allowable premium on downside risk, the market participants force the risk-neutral distribution of the index return to remain highly left-skewed even at very long horizons.

The paper is organized as follows. The next section reviews the stylized evidence on S&P 500 index returns under both the objective measure and the option-implied risk-neutral measure. Section III presents the DPL model that reconciles all the stylized evidence. Section IV calibrates the model to the S&P 500 index returns and the index options, and discusses the implications of the estimation results. Section V discusses potential model extensions and the model's applicability to other markets. Section VI concludes.

II. Review of Stylized Evidence

I review the stylized features of financial security returns based on two data sets: the times series of S&P 500 index and the European options prices on the S&P 500 index. The time series data on S&P 500 index are daily from July 3, 1962 to December 31, 2001 (9942 observations), downloaded from CRSP (Center for Research in Security Prices, University of Chicago). The S&P 500 index options data are daily quotes on out-of-the-money options from April 1999 to May 2000 across different strikes and maturities (62,950 observations). These equity index options are listed at Chicago Board of Options Exchange (CBOE). The quotes are collected by a major bank in New York city, who have also supplied the matching information on the Black-Scholes implied volatility, the spot index level, the forward price, and the interest rate corresponding to each option quote. The option maturities range from five business days to 1.8 years. Options with expiry date within a week are deleted from the sample to avoid market microstructure effects.

A. *Power Law Decay in Index Returns*

An implication of the α -stable distribution is that the tail of the distribution obeys a power law,

$$\Pr(|r| > x) = Bx^{-\alpha},$$

where r denotes a demeaned return on an asset, B is a scaling coefficient, and α is the power coefficient of the tail, often referred to as the *tail index*. Mandelbrot (1963) illustrates this power law through a double logarithm plot of probabilities $\Pr(|r| > x)$ versus x on cotton price changes. If the price change

FIG. 1. — The tail behavior of S&P 500 index returns. The plots are on S&P 500 index returns at different time horizons: daily (circle), five days (cross), 20 days (square), and 60 days (diamond). All returns are standardized by their respective sample estimates of the mean and the standard deviation. The two benchmark lines are from a standard normal distribution (dash-dotted line) and a symmetric α -stable distribution (solid line) with $\alpha = 1.9$, $\sigma = 0.7$ (scaler), and $\mu = 0$ (drift).

obeys a power law, the double logarithm plot will generate a straight line for large x , and the slope of the line becomes an estimate of the power coefficient α .

Figure 1 depicts a similar plot on the S&P 500 index log returns. I compute the log returns at different aggregation levels: daily (circle), five days (cross), 20 days (square), and 60 days (diamond). The plots for these returns are overlayed on the same figure. For ease of comparison, I standardize all returns by their respective sample estimates of the mean and the standard deviation. I also plot two benchmark lines based on a standard normal distribution (dash-dotted line) and a symmetric α -stable distribution (solid line) with $\alpha = 1.9$, $\sigma = 0.7$ (scaler), and $\mu = 0$ (drift).

The plots on the index returns approach a straight line at large values of returns, indicating the presence of power tails in the return distribution. This pattern forms a clear contrast to the curved line of the normal benchmark, the tails of which decay exponentially. The plots for the index returns at

different aggregation levels overlap one another reasonably well, indicating that the power law is fairly stable with respect to time aggregation. Both features are consistent with an α -stable distribution.

Nevertheless, comparing the data scatter plots to the α -stable distribution benchmark (solid line) reveals that the tails of index returns at very large realizations do not look as thick as the tails of the α -stable distribution, even though the benchmark plot uses a fairly large tail index at $\alpha = 1.9$. The observed data points lie between the exponential decay of a normal distribution and the power decay of an α -stable distribution.

B. Applicability of the Central Limit Theorem

In testing whether an α -stable distribution governs the asset return behavior, many empirical studies exploit the stability-under-addition property of the stable distribution. These studies estimate the tail index parameter α using data of different frequencies and analyze how the parameter estimates vary across different frequency choices. Although the results are mixed, the main finding is that the tail index estimates increase with time aggregation, a result that confirms with the traditional central limit theorem, but contradicts the implication of a pure α -stable distribution.

A simpler way of testing the stability of the return distribution is to measure the skewness and kurtosis of the asset returns under different time aggregation levels. Under the assumption of a normal distribution, both measures are zero. Under the assumption of an α -stable distribution, neither measures are well defined and hence the estimates for both should exhibit instability. Therefore, if the estimates of these moments are stable and obey the central limit theorem in converging to zero (normality) with time aggregation, the assumption of an α -stable distribution is violated.

FIG. 2. — Applicability of the central limit theorem to S&P 500 index returns. Dashed lines are estimates of skewness (left panel) and kurtosis (right panel) of the log returns on the S&P 500 index at different time aggregation levels (from one to 20 days). The solid lines are implied by the central limit theorem on iid returns with finite variance.

Figure 2 plots the skewness (left panel) and kurtosis (right panel) estimates for log returns on the S&P 500 index at different time aggregation levels. The dashed line are estimates from the data. The solid lines are inferred from the central limit theorem on iid returns with finite variance. Returns on the S&P 500 index comply well with the central limit theorem: Although the daily return distribution exhibits moderate skewness and large kurtosis, the absolute magnitudes of both statistics decline rapidly with time aggregation.

Compared to the benchmark plot for iid returns (the solid line), non-normalities in the data decay slightly slower. A slower decay can occur when the return and/or return volatility is serially correlated. Overall, the steady decline in absolute magnitudes of the skewness and kurtosis estimates supports the applicability of the central limit theorem, but contradicts the assumption of an α -stable distribution.

C. Distinct Behaviors of the Risk-Neutral Distribution

The previous subsections use the time series data to infer the properties of the return distribution under the objective measure. This subsection exploits the cross sections of the options data to analyze the return distribution under the risk-neutral measure.

Practitioners in the options market often summarize the information using the Black and Scholes (1973) implied volatility of the options. Under the normal return distribution assumption of the Black-Scholes model, this implied volatility should be a fixed number across option strike prices, or some measures of moneyness. In reality, the implied volatility often exhibits a smile or smirk pattern across moneyness as a directly result of conditional non-normality in the risk-neutral distribution of the underlying asset return. The slope of the implied volatility smirk reflects asymmetry in the risk-neutral distribution of the underlying return, and the curvature of the smirk reflects the fat-tails (leptokurtosis) of this distribution (Backus, Foresi, and Wu (1997)).

Figure 3 plots the average shapes of the option implied volatility against a standard measure of moneyness for the S&P 500 index. This moneyness measure is defined as the logarithm of the strike price over the forward, normalized by volatility and the square root of maturity. The left panel plots the nonparametrically smoothed implied volatility surface across both maturity and moneyness. The right panel plots the two-dimensional slices of the implied volatility smirk at different maturities.

At a fixed maturity level, the implied volatility smirk is highly skewed to the left, implying a highly asymmetric risk-neutral distribution for the equity index return. Across maturities, the slope of the implied volatility smirk does not flatten as maturity increases. This maturity pattern indicates that the index return distribution under the risk-neutral measure remains highly asymmetric as the time

FIG. 3. — Implied volatility smirks for S&P 500 index options. I obtain the implied volatility surface (left panel) via nonparametric smoothing of daily closing implied volatility quotes on S&P 500 index options from April 4, 1999 to May 31, 2000 (62,950 observations). The nonparametric estimation employs independent Gaussian kernels with bandwidths 0.2209 and 0.0715 along the moneyness and maturity dimension, respectively. Maturity is in years. Moneyness is defined as $d \equiv \ln(K/F)/\sigma\sqrt{\tau}$, where $\sigma = 27.4\%$ is the average of all implied volatility quotes, K is the strike price and F is the forward price. The right panel is a two-dimensional slice of the implied volatility smirks at maturities of one month (solid line), six months (dashed line), and 12 months (dash-dotted line), respectively.

horizon increases. The stability-under-addition property holds under the risk-neutral measure and up to the observable horizon of two years.¹ This stability-under-addition feature under the risk-neutral measure forms a sharp contrast to the behavior of the time series return distribution, which shows rapidly declining non-normality with increasing time aggregation.

Another distinct feature of the risk-neutral distribution for the equity index return is that it is much more skewed to the left than the return distribution under the objective measure. Figure 4 compares the nonparametrically estimated probability density function of the one-month equity index return (solid line) with the one-month conditional density inferred from the index options data (dashed line). Refer to Aït-Sahalia and Lo (1998) for the details on the nonparametric estimation of the risk-neutral density

¹More recently, Foresi and Wu (2003) find that the same maturity pattern holds for all major equity indexes in the world and for time-to-maturities up to five years.

FIG. 4. — Probability density of standardized returns on the S&P 500 index. The solid line is the nonparametrically estimated density of the standardized returns over 25 business day horizon on S&P 500 index. The dashed line is the risk-neutral conditional density computed from option prices on S&P 500 index with one month maturity. The dotted line is a standard normal benchmark.

from the options data. Figure 4 plots both densities in terms of standardized return. The dotted line represents a standard normal distribution benchmark. Compared to the normal benchmark, both the risk-neutral and the objective densities on the S&P 500 index returns are more spiked in the middle and have thicker left tails. Nevertheless, the right tail of the risk-neutral distribution is much thinner than the right tail of the objective distribution, and is even thinner than that of the normal benchmark. Thus, the risk-neutral distribution of the equity index return is much more skewed to the left than its objective counterpart.

A successful model for the equity index process should be able to reconcile the stylized evidence documented in this section. The model should generate a return distribution under the objective measure that exhibits power tails but nevertheless obeys the central limit theorem. Meanwhile, the model should also generate a risk-neutral distribution that is much more skewed to the left than its objective counterpart and that preserves stability across different time horizons.

III. A Stylized Model

In this section, I propose a stylized model that reconciles all the above stylized evidence on the equity index returns under both the objective measure and the risk-neutral measure. The model is as stylized as the evidence. The purpose of developing such a stylized model is to gain better understanding on the tail behavior of asset return innovations and to gain insights on the economic underpinnings of the distribution differences under the objective and the risk-neutral measures. The stylized model can also be used as a springboard and a key component in more comprehensive modeling endeavors.

A. The Dampened Power Law (DPL)

Let X be a one-dimensional pure jump Lévy process defined on a probability space $(\Omega, \mathcal{F}, P)$. I use X_t to capture the uncertainty of the economy and model the price of an asset S_t as an exponential affine function of X_t ,

$$S_t = S_0 \exp(\mu t + X_t - k(1)t), \quad (1)$$

where μ denotes the instantaneous drift of the asset price process and $k(1)$ is a convexity adjustment of X_t so that the term, $\exp(X_t - k(1)t)$, forms a P -martingale. This adjustment term can be derived from the cumulant exponent of X_t ,

$$k(s) \equiv \frac{1}{t} \log \mathbb{E}[e^{sX_t}], \quad s \in \mathcal{D}, \quad (2)$$

where $\mathbb{E}[\cdot]$ denotes the expectation operator under measure P and $\mathcal{D}$ denotes the subset of the real space where $k(s)$ is well defined. The cumulant exponent of a pure jump Lévy process can be computed via the Lévy-Khinchine theorem (Bertoin (1996)),

$$k(s) = \int_{\mathbb{R}^0} (e^{sx} - 1 - sh(x)) \nu(x) dx \quad (3)$$

where $\nu(x)$ is the *Lévy density* of the pure jump Lévy process X_t , which is defined on $\mathbb{R}^0$ (the real line excluding zero) and controls the arrival rate of jumps of size x . The function $h(x) : \mathbb{R}^0 \rightarrow \mathbb{R}^0$ denotes a truncation function used to analyze the jump properties around the singular point of zero jump size. It can be any function that is bounded, with compact support, and satisfies $h(x) = x$ in a neighborhood of zero ((Jacod and Shiryaev (1987))). The specification of the Lévy density controls the key feature of the model.

Definition 1 (Dampened Power Law (DPL)) *The arrival rate of jumps of size x in asset returns follows a power law, dampened by an exponential function:*

$$\nu(x) = \begin{cases} \gamma_+ e^{-\beta_+ |x|} |x|^{-\alpha-1}, & x > 0 \\ \gamma_- e^{-\beta_- |x|} |x|^{-\alpha-1}, & x < 0 \end{cases} \quad (4)$$

with the parameters $\alpha \in (0, 2]$, $\beta_{\pm}, \gamma_{\pm} \in \mathbb{R}^+$.

By setting $\beta_{\pm} = 0$ and hence without exponential dampening, the Lévy density uniquely determines an α -stable Lévy motion that generates the α -stable distribution proposed by Mandelbrot (1963) and Fama (1965). The arrival rate of jumps of size x decays in power law. The difference in γ_+ and γ_-

determines the asymmetry of the α -stable distribution.² As a special example, Carr and Wu (2003) set $\gamma_+ = 0$ so that they only allow negative jumps in their model.

With strictly positive dampening $\beta_{\pm} > 0$, the exponential functions $e^{-\beta_+|x|}$ and $e^{-\beta_-|x|}$ in equation (4) dampen the Lévy density so that the arrival rate of jumps decays faster as the absolute jump size $|x|$ increases. I label $\beta_{\pm}$ as the *dampening coefficients* and say that the asset return innovation X_t obeys the dampened power law (DPL).

Carr, Geman, Madan, and Yor (2002) consider a similar specification for the Lévy measure, but with the constraint of $\gamma_+ = \gamma_-$. They regard the specification as an extension to the variance gamma model of Madan, Carr, and Chang (1998) and Madan and Seneta (1990), where $\alpha = 0$ and thus the return distribution does not have a power component. If we set $\alpha = -1$, the specification in equation (4) captures a double-exponential specification as in Kou (2002). However, my focus in this paper is on models with a power decay and hence a strictly positive α .

The exponential dampening dramatically alters the fundamental properties of the return innovation X_t . Without dampening, X_t follows an α -stable distribution and only moments of order less than α are well-defined. Given $\alpha < 2$, the variance of the return is not finite, and hence the classic central limit theorem does not apply. With strictly positive dampening ($\beta_{\pm} > 0$), the following proposition states that the moments of X_t of all finite orders are finite.

²In principle, the power α can also be different for the two sides of the distribution, but we follow convention in introducing asymmetry only through $\gamma_{\pm}$. Refer to Janicki and Weron (1994), Samorodnitsky and Taqqu (1994), and Zolotarev (1986) for details on α -stable distributions.

Proposition 1 *Given the Lévy density in equation (4), with strictly positive dampening ($\beta_{\pm} > 0$) and with $\alpha \neq 1$, the cumulant exponent of X_1 is*

$$k(s) = \Gamma(-\alpha)\gamma_+[(\beta_+ - s)^\alpha - \beta_+^\alpha] + \Gamma(-\alpha)\gamma_-[(\beta_- + s)^\alpha - \beta_-^\alpha] + sC(h), \quad (5)$$

where $C(h)$ is an immaterial constant that depends on the exact form of the truncation function $h(x)$ but will be eventually cancelled out with the convexity adjustment term in the asset price specification.

The j -th cumulant is given by

$$\kappa_1 \equiv \left. \frac{\partial k(s)}{\partial s} \right|_{s=0} = \Gamma(1-\alpha) [\gamma_+ (\beta_+)^{\alpha-1} - \gamma_- (\beta_-)^{\alpha-1}] + C(h), \quad (6)$$

$$\kappa_j \equiv \left. \frac{\partial^j k(s)}{\partial s^j} \right|_{s=0} = \Gamma(j-\alpha) [\gamma_+ (\beta_+)^{\alpha-j} + (-1)^j \gamma_- (\beta_-)^{\alpha-j}], j = 2, 3, \dots, \quad (7)$$

which is finite for all j as long as $\beta_{\pm} > 0$. When either $\beta_+ = 0$ or $\beta_- = 0$, only moments of order less than $\alpha \leq 2$ are finite.

I leave the proof in Appendix A. The cumulant exponent takes a different form for the special case of $\alpha = 1$, the results of which is in Appendix B. The other special case is when $\alpha = 0$, i.e., the variance-gamma model, which I refer the interested readers to Madan, Carr, and Chang (1998) and Madan and Seneta (1990) for details. For ease of exposition, I will base the discussions in the paper on the general case with $\alpha \neq 1$.

The return innovation X_t has finite moments of all orders as long as the dampening coefficients on both sides of the distribution are strictly positive. Without dampening on either side, the variance of the asset return does not exist and hence the central limit theorem does not apply.

For asset pricing, we are concerned not only with the finiteness of moments of the asset return, but also with the finiteness of moments of the asset price. For example, if the conditional mean (first moment) of the asset price were not finite under the risk-neutral measure, there would not exist a martingale measure with a finite interest rate. The no-arbitrage condition might then be violated, a concern originally raised by Merton (1976) on the applicability of α -stable distributions in modeling asset returns. Under the DPL specification, the dampening coefficients (β_+, β_-) also determine the existence of price moments.

Proposition 2 *With $\gamma_{\pm}, \gamma_{\pm} > 0$, the cumulant exponent of X is well defined on $s \in (-\beta_-, \beta_+)$.*

By the definition of the cumulant function, this means that the conditional price moments are finite within the orders of $(-\beta_-, \beta_+)$. Thus, under the specification in equation (1), for the convexity adjustment term $k(1)$ to be finite and with $\gamma_+ > 0$, the dampening coefficient on the positive jumps, β_+ , must be no less than one. The proof for this proposition follows the proof of Proposition 1 in Appendices A and B.

B. The Market Price of Jump Risk

Consistent with the separate parameterization on the arrival rate of negative and positive jumps, I also allow market participants to have different risk attitudes toward positive and negative jumps. For example, for a security with an aggregate long position in the market, such as an equity index, downside and upside jumps generate quite different impacts on people's wealth. Thus, it is very likely that the market treats the downside jumps as "hazards" and upside jumps as "potentials" and charge different premiums on jumps of different directions. In contrast, for a process underlying a net zero position,

such as an exchange rate process, jumps of both directions are more likely to be treated equally. My separate treatment of downside and upside risks allows the data to determine whether or not the market discriminates asset price movement of different directions.

Formally, corresponding to an instantaneous interest rate r , I define a new measure Q that is absolutely continuous with respect to the objective measure P . Under this measure Q , asset prices discounted by the bank account defined on r become martingales. No arbitrage guarantees the existence of at least one such measure, often referred to as the risk-neutral measure. I propose that the following extended exponential martingale defines the measure change from P to Q ,

$$\left. \frac{dQ}{dP} \right|_t \equiv \exp(-\lambda_+ X_t^+ - tk_+(-\lambda_+)) \exp(-\lambda_- X_t^- - tk_-(-\lambda_-)), \quad (8)$$

where X^+ and X^- are independent processes consisting only of the positive and negative jumps of X , respectively, with $X = X^+ + X^-$. Accordingly, k_+ and k_- are the cumulant exponents of X^+ and X^- , respectively, with $k = k_+ + k_-$. My extension to the standard exponential martingale lies in the different parameterizations λ_+ and λ_- for positive and negative jumps, respectively. The literature refers to λ as the *market price of risk*. Under my extension, λ_+ is the market price of upside jump risk and λ_- is the market price of downside jump risk.

Applications of exponential martingales for measure changes, normally without the separate treatment for positive and negative jumps, have been considered in Carr and Wu (2004), Madan and Milne (1991), Gerber and Shiu (1994), Eberlein and Keller (1995), and Kallsen and Shiryaev (2002). Kallsen (2000) considers its link to exponential utility maximization. Miyahara (1999) and Chan (1999) consider its link to the relative entropy minimization.

Given the specification of the Lévy density $v(x)$ of X_t in equation (4) under measure P and the measure change defined in equation (8), the Lévy density of X_t under measure Q becomes³

$$v^Q(x) = \begin{cases} e^{-\lambda_+ x} v(x) = \gamma_+ e^{-(\beta_+ + \lambda_+) |x|} |x|^{-\alpha-1}, & x > 0 \\ e^{-\lambda_- x} v(x) = \gamma_- e^{-(\beta_- - \lambda_-) |x|} |x|^{-\alpha-1}, & x < 0 \end{cases} \quad (9)$$

If I further define $\beta_+^Q \equiv \beta_+ + \lambda_+$ and $\beta_-^Q \equiv \beta_- - \lambda_-$, it becomes obvious that X_t also obeys an exponentially dampened power law under the risk-neutral measure Q . The dampening coefficients for positive and negative jumps under the risk-neutral measure are β_+^Q and β_-^Q , respectively.

By analogy to Proposition 1, with $\alpha \neq 1$, the cumulant exponent of X_t under the new measure Q is,

$$k^Q(s) = \Gamma(-\alpha) \gamma_+ \left[\left(\beta_+^Q - s \right)^\alpha - \left(\beta_+^Q \right)^\alpha \right] + \Gamma(-\alpha) \gamma_- \left[\left(\beta_-^Q + s \right)^\alpha - \left(\beta_-^Q \right)^\alpha \right] + s C^Q(h), \quad (10)$$

which is finite for $s \in (-\beta_-^Q, \beta_+^Q)$. Furthermore, the j -th cumulant of X_1 under measure Q is

$$\kappa_1 = \Gamma(1-\alpha) \left[\gamma_+ \left(\beta_+^Q \right)^{\alpha-1} - \gamma_- \left(\beta_-^Q \right)^{\alpha-1} \right] + C^Q(h), \quad (11)$$

$$\kappa_j = \Gamma(j-\alpha) \left[\gamma_+ \left(\beta_+^Q \right)^{\alpha-j} + (-1)^j \gamma_- \left(\beta_-^Q \right)^{\alpha-j} \right], j = 2, 3, \dots \quad (12)$$

which is finite for all $j = 1, 2, \dots$ as long as $\beta_\pm^Q \neq 0$. When either $\beta_+^Q = \beta_+ + \lambda_+ = 0$ or $\beta_-^Q = \beta_- - \lambda_- = 0$, only moments of order less than $\alpha \leq 2$ are finite.

Under this risk-neutral measure Q , the asset price S_t becomes

$$S_t = S_0 e^{(r-q)t + X_t - t k^Q(1)}, \quad (13)$$

³Refer to Küchler and Sørensen (1997) for measure changes under exponential martingales.

where $k^Q(1)$ is given in (10) and q is the dividend yield. No arbitrage dictates that the instantaneous drift is $(r - q)$ under the risk-neutral measure. The market risk premium on the asset return is given by

$$\mu - (r - q) = k(1) - k^Q(1).$$

The exponential martingale has an asymmetric flavor in its definition of market price of risk even if $\lambda_+ = \lambda_- = \lambda$. In particular, a positive market price of risk λ fattens the left tail of the asset return (negative jumps) but thins the right tail (positive jumps) of the asset return under the risk-neutral measure. This asymmetry generates the difference between $k(1)$ and $k^Q(1)$ and hence the risk premium in return. Thus, starting at a symmetric distribution under the objective measure P , the risk-neutral density of the return distribution becomes skewed to the left when the market price of risk λ is positive and of the same magnitude on jumps of both directions.

Since $\beta_-^Q = \beta_- - \lambda_-$ needs to be nonnegative for the Lévy density to be well defined, to exclude arbitrage, the market price on downside jumps is bounded from above at $\lambda_- \leq \beta_- \leq 0$. That is, even if the market is extremely averse to downward jumps, under no arbitrage, the maximum premium that can be charged on the downward jumps is $\lambda_- = \beta_-$.

Remark 1 (Unique Feature of S&P 500 Index Options) *For S&P 500 index options, if $\lambda_- = \beta_-$ and the market charges the maximum premium allowable by no arbitrage on downside index jumps, the left tail of the risk-neutral distribution of the index return follows a power law with no dampening.*

With $\lambda_- = \beta_-$, return variance and higher moments are infinite. Hence, the central limit theorem does not apply to the asset return under measure Q .

Therefore, by modeling asset returns with the exponentially dampened power law, I can reconcile all the stylized findings documented in Section II. With exponential dampening, asset returns can both have power tails and obey the central limit theorem in converging to normality with time aggregation. Furthermore, when the market charges the maximally allowable premium on downside index movement, the left tail of the risk-neutral return on the index is no longer dampened and hence the central limit theorem no longer applies, consistent with the observation from the index options market.

IV. Calibration Exercises

To gain further insights on the model and gauge the market attitudes toward downside and upside movements in the equity index, I calibrate the DPL model to both the time series of the S&P 500 index returns and the cross-section of the option prices on S&P 500 index. The model parameters vector is given by $\Theta = [\mu, \alpha, \gamma_{\pm}, \beta_{\pm}, \beta_{\pm}^Q]^{\top}$. I calibrate two versions of the model, one being unconstrained, the other with the constraint $\beta_{-}^Q = 0$ and hence $\lambda_{-} = \beta_{-}$, under the null hypothesis that the market charges the maximally allowable premium on downward index movement. Performance comparisons between the two versions of the model shed light on whether the market charges the maximum premium on the downside index movement.

A. Data and Estimation

The data sources for the equity index returns and the equity index options are described in Section II. For the time series data, to increase the stability of the numerical algorithm, I calibrate the models to standardized log returns, i.e. returns that are demeaned and normalized by its sample standard deviation.

tion. Furthermore, due to the telescopic property of the log returns, an arithmetic sample average would present a noisy estimate of the mean return that only depends upon the first and the last observation. Instead, I estimate the mean return by regressing log price on time t ,

$$\ln S_t = a + bt + e,$$

where $t = [1 : T]/252$, with T being the number of daily observations, and the estimate for b is an estimate for the mean annualized log return. Based on other model parameter estimates and the model specification in (1), the estimate for the instantaneous drift of the index is given by

$$\mu = \widehat{b} + \kappa_1 - k(1),$$

where the κ_1 is the first cumulant of X_1 and $k(1)$ is the cumulant exponent of X . Recall that X defines the uncertainty of the economy and is described by the DPL Lévy density in equation (4).

To facilitate estimation, I also normalize the option prices as the forward option price in percentages of the forward underlying price,

$$p(k, \tau) = 100 \times \frac{P(k, \tau)e^{r\tau}}{F},$$

where $P(k, \tau)$ denotes the out-of-the-money option midquote at moneyness k and maturity τ . The moneyness in this case is defined as $k \equiv \ln K/F$. Under the Lévy assumption, this normalized option price at each fixed moneyness and maturity should be identical across different dates. Thus, I can estimate the mean value and variance of the normalized option price at each moneyness and maturity via nonparametric regression.

I use the fast Fourier transform (FFT) method of Carr and Madan (1999) to compute model price for the options based on the characteristic function of the log return. Since this FFT algorithm generates option prices at fixed moneyness with equal intervals at each maturity, options at observed maturities are used for the estimation. But at each maturity, I sample the options data with a fixed moneyness interval of $\Delta k = 0.03068$, within the moneyness range $k = \ln K/F = (-0.3988, 0.1841)$. This moneyness range excludes approximately 16 percent deep out-of-the-money options (approximately eight percent calls and eight percent puts) which I deem as too illiquid to contain useful information. I apply linear interpolation to obtain the option prices at the fixed moneyness grids, resulting in a maximum of 20 strike points at each maturity. For the interpolation to work with sufficient precision, I require that there be at least five data points at each date and maturity. I also refrain from extrapolating by only retaining option prices at fixed moneyness intervals that are within the data range. Visual inspection indicates that at each date and maturity, the quotes are so close to each other along the moneyness line that interpolation can be done with little error, irrespective of the interpolation methods. In total, the procedure generates 35,038 option sample data points used for estimation.

With the above data set, I estimate the models using a maximum likelihood method. Under the Lévy specification, stock returns are independently and identically distributed under both the objective measure P and the risk-neutral measure Q . I exploit this property to expedite the likelihood calculation. First, given the cumulant exponent expressions in (5) and (10), the characteristic functions of the stock returns over horizon t , $s_t = \ln(S_t/S_0)$, are given by,

$$\begin{aligned}\varphi(u) &\equiv \mathbb{E} [e^{ius_t}] = e^{iu(\mu - k(1))t + tk(iu)}, \\ \varphi^Q(u) &\equiv \mathbb{E}^Q [e^{ius_t}] = e^{iu(r - q - k^Q(1))t + tk^Q(iu)},\end{aligned}\tag{14}$$

under measures P and Q , respectively. Second, given the characteristic function of the log returns under measure P , I apply the fast Fourier transform (FFT) to efficiently compute the probability density at a fine grid of return levels. The log likelihood value of the time series return data can thus be readily computed from these densities.

The likelihood for the option prices is computed by assuming that the option pricing errors are normally distributed. Given the Lévy specification, the normalized option price of the index, $p(k, \tau)$, should be the same across different days at fixed moneyness and maturity levels. Thus, a mean estimate of the option price at each moneyness and maturity reflects its “true” value and the variance estimate reflects the variance of the pricing error. Assuming that the pricing errors are independently, normally distributed with distinct variance at different moneyness and maturity, I construct the log likelihood function based on the normalized option price as

$$l(k, \tau) = -\frac{1}{2} \ln \left(2\pi \widehat{V}(k, \tau) \right) - \frac{(p(k, \tau) - p(k, \tau; \Theta))^2}{2\widehat{V}(k, \tau)},$$

where $p(k, \tau)$ is the normalized option price at moneyness k and maturity τ , $p(k, \tau; \Theta)$ is the corresponding model value with parameter vector Θ , and $\widehat{V}(k, \tau)$ denotes the variance estimate of the pricing error at moneyness k and maturity τ . Under the Lévy specification and given the Fourier transform of the risk-neutral return $\varphi^Q(u)$ in equation (14), the model value $p(k, \tau; \Theta)$ can be computed via the fast Fourier transform method by setting $S = F = 100$ and $r = q = 0$. Finally, since option quotes are observed at varying moneyness and maturities, I use nonparametric regression to estimate the sample variance, $\widehat{V}(k, \tau)$, of the normalized option quotes at each fixed moneyness and maturity level. I apply independent Gaussian kernels for the nonparametric regression, with bandwidths at 0.1386 and 0.2862 along the moneyness and maturity dimension, respectively.

The aggregate likelihood function ($\mathcal{L}$) is then constructed as a summation of the log likelihood from the time series returns and the log likelihood from the cross-section of options. The model parameters are estimated by maximizing the aggregate likelihood value.

B. Model Parameter Estimates

Table 1 presents the model parameter estimates, together with their standard errors and p -values. Given the extremely large sample used for the estimation, the standard errors for most model parameters are very small, so are their p -values. Panel A of Table 1 contains the estimates for the unconstrained model. The tail index α is 1.4892, close to literature estimates on pure α -stable models without dampening. The scaling coefficients $\gamma_{\pm}$ control the asymmetry of the distribution in the absence of exponential dampening. The scale estimate on positive jumps, $\gamma_+ = 0.0024$, is more than ten times smaller than the scale estimate on negative jumps, $\gamma_- = 0.0315$. Without dampening or with symmetric dampening, this different scaling generates negative skewness in the return distribution.

Now we look at the dampening coefficients $\beta_{\pm}$ on both tails under the objective measure. These dampening coefficients influence the tail behavior of the return distribution under the objective measure. The dampening coefficient on the right tail is fairly moderate at $\beta_+ = 1.0015$, just barely enough to guarantee the existence of the first price moment. The dampening on the left tail is much stronger at $\beta_- = 12.9788$. Since the scaling coefficient on positive jumps γ_+ is much smaller than the scaling coefficient on negative jumps γ_- , the lighter dampening counteracts with the smaller scaling on the right tail to make it similar to the left tail, which has a larger scaling coefficient, but is also dampened more heavily. The net result of the interactions between dampening and scaling is a relatively symmetric return distribution under the objective measure (see the solid line in Figure 4).

The estimates for the risk-neutral dampening coefficients look dramatically different from their objective-measure counterparts. Under the risk-neutral measure, positive jumps are dampened much more heavily than under the objective measure (5.2306 for β_+^Q versus 1.0015 for β_+). This heavy dampening, combined with the small scaling (γ_+), makes the right tail very thin under the risk-neutral measure. In contrast, the dampening on the left tail of the risk-neutral distribution is negligible, with the estimate for β_-^Q very close to zero at 0.0067. This is dramatically different from the heavy dampening under the objective measure ($\beta_- = 12.9788$). This negligible dampening, together with the large scaling parameter γ_- generates a very fat left tail for the risk-neutral return distribution, supporting the evidence in Figure 4 (the dashed line).

The estimate for β_-^Q is the only estimate that has a large p -value (0.7656) and hence is not significantly different from zero. With β_-^Q at zero and therefore no dampening on the left tail, the index return exhibits infinite variance under the risk-neutral measure. The classic central limit theorem no longer applies, and the return non-normality persists as option maturity increases. Thus, we achieve stability under time aggregation on the model-generated implied volatility smirk across different maturities, in line with the observation in Figure 3.

The differences between the dampening coefficients under the risk-neutral measure and the objective measure capture the market price of risk. The market price of upside jump risk is $\lambda_+ = \beta_+^Q - \beta_+ = 5.2306 - 1.0015 = 4.2291$. The positive λ_+ estimate implies a thinner right tail under the risk-neutral measure than under the objective measure. It represents a discounting of the positive index movement to compensate for uncertainty.

The market price of downward jump risk is $\lambda_- = \beta_- - \beta_-^Q = 12.9788 - 0.0067 = 12.9721$. The positive λ_- estimate implies a thicker left tail under the risk-neutral measure than under the objective

measure. It represents a premium charged against the downward index movement. The fact that both estimates are positive indicates that market participants treat unanticipated shocks in both directions as risks and charge a risk premium for both directions of shocks. Furthermore, the different magnitudes of λ_+ and λ_- indicate that the market's risk attitudes toward the two directions of index movements are different. The market charges a much higher price (12.9721) for downward index movements than for upward movements (4.2291). Indeed, the premium charged on the downward index movement approaches the maximum value allowable by no arbitrage because the estimate for β_-^Q is no longer significantly different from zero.

Panel B of Table 1 reports the parameter estimates of a restricted version of the DPL model where the market price of downside risk is set to the maximum that is allowable by no arbitrage: $\lambda_- = \beta_-$ and hence $\beta_-^Q = 0$. A significant degeneration of model performance would reject this hypothesis. Compared to the unrestricted model in Panel A, the likelihood value of this restricted version is not much smaller. A likelihood ratio test between the two models, $\chi^2(1) = 2(\mathcal{L}_A - \mathcal{L}_B)$, generates a p -value of 0.3076, implying that the unrestricted version (A) of the model does not significantly outperform the restricted version (B). The estimates for other model parameters are also very similar under the two models. Therefore, the null hypothesis $\beta_-^Q = 0$ is in compliance with the data: The market charges the maximally allowable premium on downside index movements.

V. Further Applications and Extensions

I have reviewed the stylized evidence and calibrated the models using data on the S&P 500 index. In this section, I show that the DPL specification is equally applicable to the currency market. Furthermore,

by focusing on the tail behavior of asset returns, I have thus far ignored evidence on stochastic volatility and the presence of a diffusive component. I address such model extensions in this section.

A. Applicability of the DPL to the Currency Market

Power tails are not a unique feature of the equity market. Similar tail behavior has also been observed for currency returns (Calvet and Fisher (2002)). Such evidence suggests that the DPL specification could also be applicable to the currency market.

Currency options exhibit different behaviors from that of the equity index options. For comparison, Figure 5 plots the nonparametrically smoothed implied volatility surface and its two-dimensional slices on European options on Deutsche mark. The options are listed at the Philadelphia Stock Exchange (PHLX) and are downloaded from WRDS (Wharton Research Data Services). The options are daily closing quotes from September 2, 1987 to December 19, 1997. The data set also contains the corresponding spot price of the currency, along with the strike and maturity information. Domestic and foreign interest rates are based on the corresponding LIBOR rates, downloaded from Datastream. I check the no-arbitrage bounds and compute the Black-Scholes implied volatility for each option quote. The cleaned-up data set has 12,465 option quotes. The smoothed implied volatility surface in Figure 5 is from this cleaned data set.

Compared to the average implied volatility surface on the equity index in Figure 3, Figure 5 shows two sharp differences for the implied volatility surface on currency options. First, in contrast to the highly skewed feature of the implied volatility smirk for the equity index options, the implied volatility smile for the currency options is relatively symmetric. This symmetric smile implies a relatively symmetric risk-neutral distribution for the currency returns. Second, although the implied volatility smirk

FIG. 5. — Implied volatility smiles for European options on Deutsche mark. I obtain the implied volatility surface in the left panel via nonparametric smoothing of daily closing implied volatilities on European options on Deutsche mark from September 2, 1987 to December 19, 1997 (12,465 observations). Maturity is in years. Moneyness is defined as $d \equiv \ln(K/F)/\sigma\sqrt{\tau}$, where $\sigma = 11.58$ percent is the average of all implied volatility quotes, K is the strike price and F is the forward price. The right panel is a two-dimensional slice of the implied volatility smiles at maturities of one month (solid line), three months (dashed line), and six months (dash-dotted line), respectively.

on the index options does not flatten as option maturity increases, the implied volatility smile on the currency options flattens steadily with increasing maturity. Therefore, the conditional non-normality on the risk-neutral distribution of the currency return declines steadily as the conditioning horizon increases.

The DPL model can accommodate both differences by a judicious choice of the market prices of downside and upside risks ($\lambda_{\pm}$). The relatively symmetric nature of the currency return distribution under both the objective measure and the risk-neutral measure implies similar dampening and scaling coefficients for both upward and downward currency movements under both measures. Furthermore, as long as the dampening coefficients are strictly positive and similar for both tails, the conditional return non-normality will decline with the conditioning horizon, as implied by the central limit theorem. A line for future research is to calibrate the DPL model to the currency time series returns and the

currency option prices and to investigate the differences between the parameter estimates from the currency market and those from the equity market.

B. The Presence of a Diffusion Component and Stochastic Volatility

The proposition of a pure jump Lévy process, the DPL, is consistent with this paper's focus on the tails of the return distribution. Naturally, the DPL model should not be regarded as the final answer to modeling financial asset returns, but rather as an organic component of a more sophisticated model that may also include a diffusion component and stochastic volatility.

Recent empirical studies on the S&P 500 index returns and index options have come to three major findings. First, the index return process contains both a diffusion component and a jump component. Second, return volatilities are stochastic and are correlated with the return innovation, i.e. the so-called “leverage effect” (Black (1976)). Finally, stochastic volatility can come from both diffusions and jumps (Bates (2000), Huang and Wu (2004), and Pan (2002)). As an illustration, I propose an extended model structure that accommodates all the above pieces of evidence, with DPL being the centerpiece of the jump component specification.

Under the objective measure P , I propose the following process for the asset price movement,

$$S_t = S_0 \exp \left(\mu t + \sigma W_{T_t^d} - \frac{1}{2} \sigma^2 T_t^d + X_{T_t^j} - k(1) T_t^j \right), \quad (15)$$

where μ is the instantaneous drift of the asset price, σ is a positive constant, W_t is a standard Brownian motion, X_t denotes the DPL jump component, as specified by the Lévy density in equation (4). The vector $T_t \equiv [T_t^d, T_t^j]^\top$ denotes a stochastic time change applied to the two Lévy components W_t and X_t .

By definition, the time change T_t is an increasing, right-continuous process with left limits satisfying the usual regularity conditions.

I further restrict T_t to be continuous and differentiable with respect to t . In particular, let

$$v(t) \equiv \left[v^d(t), v^j(t) \right]^\top = \partial T_t / \partial t. \quad (16)$$

Then, $v^d(t)$ is proportional to the instantaneous variance of the diffusion component, and $v^j(t)$ is proportional to the arrival rate of the jump component. Following Carr and Wu (2004), I label $v(t)$ as the instantaneous activity rate and let the two activity rates follow separate stochastic processes. The following specifications for the activity rate processes represent a reasonable and parsimonious choice,

$$\begin{aligned} dv^d(t) &= \kappa^d (1 - v^d(t)) dt + \sigma_v^d \sqrt{v^d(t)} dZ_t, \\ dv^j(t) &= \kappa^j (1 - v^j(t)) dt - \sigma_v^j dX_{T_t^j}^-, \end{aligned}$$

where Z denotes another standard Brownian motion, correlated with W by $\rho dt = \mathbb{E}[dZ_t dW_t]$, X_t^- denotes the negative jumps in X_t .

This specification tightly knits the three key elements of the asset price behavior into one framework, with W_t denoting the diffusion component, X_t the jump component, and $v(t)$ the two sources of stochastic volatility. The leverage effect is incorporated via both jumps and diffusion. Leverage via diffusion is captured by a negative correlation ρ between the two Brownian motions W and Z . Leverage via jumps is captured by the synchronous movement of the negative jumps in returns and positive jumps in volatility. The notation $-X_{T_t^j}^-$ implies that whenever the return innovation X_t jumps downward, the volatility innovation jumps upward. An analogous specification can be assumed under the risk-neutral

measure. Under this specification, I can derive the characteristic function of the asset return following the method proposed in Carr and Wu (2004). A line of future research is to investigate the empirical estimation and performance of such stochastic volatility models in capturing the behaviors of different financial markets.

VI. Concluding Remarks

I propose a stylized model that can reconcile a series of seemingly conflicting findings on financial security returns and option prices. The model is based on a pure jump Lévy process, wherein the arrival rate of jumps obeys a power law dampened by an exponential function. The power law specification accommodates the historical evidence on α -stable tails observed on the returns of many financial assets. The exponential dampening generates finite return variance such that the return non-normality declines with time aggregation as a result of the classic central limit theorem. This property answers the more recent criticism and empirical evidence against the traditional α -stable specification. Furthermore, by applying an extended exponential martingale for measure change, I allow the risk premiums for upside and downside asset price movements to be different. When the risk premium on the downside movement approaches the maximum value allowable by no arbitrage, the dampening on the left tail disappears under the risk-neutral measure. Return variance becomes infinite under such a measure and the classic central limit theorem no longer applies, thus complying with the evidence on the equity index options.

I calibrate the model to S&P 500 index returns and index option prices. The model parameter estimates confirm my conjecture that the market participants' risk attitudes toward upside and down-

side index movements are quite different. The market participants only charge a moderate premium for upward index movements, but they charge the maximally allowable premium on downward index movements.

As examples for further applications and extensions, I show how the model can also be applied to the currency market. I also show how this stylized model can be extended to accommodate a diffusion component, separate sources of stochastic volatility, and the leverage effect. Further research can be devoted to investigate the empirical performance of this extended model in capturing the behavior of returns on different financial assets.

Appendix A. Proof of Propositions 1 and 2

I apply the Lévy-Khinchine theorem to the DPL Lévy density,

$$\begin{aligned} k(s) &= \int_0^\infty (e^{sx} - 1 - sh(x)) \gamma_+ e^{-\beta_+ x} x^{-\alpha-1} dx + \int_{-\infty}^0 (e^{sx} - 1 - sh(x)) \gamma_- e^{\beta_- x} |x|^{-\alpha-1} dx \\ &= k_+(s) + k_-(s). \end{aligned}$$

To perform the integration, I need to choose a truncation function. It is convenient to choose $h(x) = xI_{|x|<1}$, which satisfies all the necessary properties for a truncation function.

The cumulant exponent for positive jumps is,

$$k_+(s) = \gamma_+ \int_0^\infty (e^{sx} - 1 - sxI_{|x|<1}) e^{-\beta_+ x} x^{-\alpha-1} dx = I_1^+ + I_2^+,$$

with

$$\begin{aligned} I_1^+ &= \gamma_+ \int_0^\infty (e^{sx} - 1 - sx) e^{-\beta_+ x} x^{-\alpha-1} dx, \\ I_2^+ &= \gamma_+ \int_0^\infty (sx - sxI_{|x|<1}) e^{-\beta_+ x} x^{-\alpha-1} dx. \end{aligned}$$

For I_1^+ , I first Taylor expand the exponential function and then integrate term by term,

$$I_1^+ = \gamma_+ \sum_{m=2}^\infty \frac{1}{m!} \int_0^\infty (sx)^m e^{-\beta_+ x} x^{-\alpha-1} dx = \gamma_+ \sum_{m=2}^\infty \frac{1}{m!} s^m \int_0^\infty (x)^{m-\alpha-1} e^{-\beta_+ x} dx \quad (\text{A1})$$

$$= \gamma_+ \sum_{m=2}^\infty \frac{1}{m!} s^m \beta_+^{-(m-\alpha)} \Gamma(m-\alpha) = \gamma_+ \beta_+^\alpha \sum_{m=2}^\infty \frac{1}{m!} \left(\frac{s}{\beta_+} \right)^m \Gamma(m-\alpha). \quad (\text{A2})$$

From (A1) to (A2), I apply the Gamma function,

$$\Gamma(t) = \int_0^\infty e^{-x} x^{t-1} dx, \quad t > 0.$$

The series in (A2) is a real-valued, convergent series as long as $s < \beta_+$. Assume that this is the case and that $\alpha \neq 1$, I can consolidate the series expansion into the following,

$$I_1^+ = \gamma_+ \beta_+^\alpha \left(\frac{s}{\beta_+} \alpha - 1 + \left(1 - \frac{s}{\beta_+} \right)^\alpha \right) \Gamma(-\alpha) = \gamma_+ \Gamma(-\alpha) ((\beta_+ - s)^\alpha - \beta_+^\alpha + s\alpha\beta_+^{\alpha-1}).$$

For the second part, I have

$$I_2^+ = \gamma_+ \int_1^\infty s x^{-\alpha} e^{-\beta_+ x} dx = s \gamma_+ \beta_+^{\alpha-1} \Gamma(1-\alpha, \beta_+).$$

Combining I_1^+ and I_2^+ generates

$$\begin{aligned} k_+(s) &= I_1^+ + I_2^+ \\ &= \gamma_+ \Gamma(-\alpha) ((\beta_+ - s)^\alpha - \beta_+^\alpha + s\alpha\beta_+^{\alpha-1}) + s \gamma_+ \beta_+^{\alpha-1} \Gamma(1-\alpha, \beta_+) \\ &= \gamma_+ \Gamma(-\alpha) ((\beta_+ - s)^\alpha - \beta_+^\alpha) + s C_+, \end{aligned}$$

with

$$C_+ = \gamma_+ \beta_+^{\alpha-1} [\Gamma(-\alpha) \alpha + \Gamma(1-\alpha, \beta_+)].$$

The convexity adjustment term for the upside jump is,

$$k_+(1) = \gamma_+ \Gamma(-\alpha) ((\beta_+ - 1)^\alpha - \beta_+^\alpha) + C_+.$$

It is obvious that the linear drift term C_+ will be cancelled out in the convexity-adjusted jump process $X_t^+ - tk(1)$.

Hence, C_+ is immaterial for my analysis and model estimation.

Now I turn to the cumulant exponent for negative jumps

$$\begin{aligned}
k_-(s) &= \gamma_- \int_{-\infty}^0 (e^{sx} - 1 - sxI_{|x|<1}) e^{-\beta_-|x|} |x|^{-\alpha-1} dx \\
&= \gamma_- \int_0^{\infty} (e^{-sx} - 1 + sxI_{|x|<1}) e^{-\beta_-x} x^{-\alpha-1} dx = I_1^- + I_2^-.
\end{aligned}$$

The two integrals can be derived analogously,

$$\begin{aligned}
I_1^- &= \gamma_- \int_0^{\infty} (e^{-sx} - 1 + sx) e^{-\beta_-x} x^{-\alpha-1} dx = \gamma_- \sum_{m=2}^{\infty} \frac{1}{m!} \int_0^{\infty} (-sx)^m e^{-\beta_-x} x^{-\alpha-1} dx \\
&= \gamma_- \sum_{m=2}^{\infty} \frac{1}{m!} (-s)^m \int_0^{\infty} (x)^{m-\alpha-1} e^{-\beta_-x} dx = \gamma_- \beta_-^{\alpha} \sum_{m=2}^{\infty} \frac{1}{m!} \left(-\frac{s}{\beta_-}\right)^m \Gamma(m-\alpha). \quad (\text{A3})
\end{aligned}$$

For the series to be convergent and real valued, I need $s > -\beta_-$. Assume that this is true and that $\alpha \neq 1$, I have

$$I_1^- = \gamma_- \Gamma(-\alpha) ((\beta_- + s)^{\alpha} - \beta_-^{\alpha} - s\beta_-^{\alpha-1}\alpha).$$

The second integral is

$$I_2^- = -\gamma_- \int_1^{\infty} sx^{-\alpha} e^{-\beta_-x} dx = -s\gamma_- \beta_-^{\alpha-1} \Gamma(1-\alpha, \beta_-).$$

Therefore,

$$k_-(s) = \gamma_- \Gamma(-\alpha) ((\beta_- + s)^{\alpha} - \beta_-^{\alpha}) + sC_-,$$

with the immaterial linear term

$$C_- = -\gamma_- \beta_-^{\alpha-1} [\Gamma(-\alpha)\alpha + \Gamma(1-\alpha, \beta_-)].$$

Combining the cumulants for negative and positive jumps together, I have

$$k(s) = \gamma_+ \Gamma(-\alpha) ((\beta_+ + s)^{\alpha} - \beta_+^{\alpha}) + \gamma_- \Gamma(-\alpha) ((\beta_- + s)^{\alpha} - \beta_-^{\alpha}) + sC, \quad (\text{A4})$$

with $C = C_+ + C_-$, under the assumption that $\alpha \neq 1$ and $s < \beta_+, s > -\beta_-$, or $s \in (-\beta_-, \beta_+)$. Equation (2) proves equation (5) in Proposition 1. The assumptions on s are necessary for the cumulant exponent to be convergent. Hence, Proposition 2 is also proved.

The cumulants of X_1 can be obtained by progressively evaluating the derivative $\kappa_j = \left. \frac{\partial k(s)}{\partial s} \right|_{s=0}$:

$$\begin{aligned} \kappa_1 &= \Gamma(-\alpha) \gamma_+ \left[-\alpha (\beta_+ - s)^{\alpha-1} \right] + \Gamma(-\alpha) \gamma_- \left[\alpha (\beta_- + s)^{\alpha-1} \right] + C \\ &= \Gamma(1-\alpha) \left[\gamma_+ (\beta_+)^{\alpha-1} - \gamma_- (\beta_-)^{\alpha-1} \right] + C, \\ \kappa_j &= \Gamma(j-\alpha) \left[\gamma_+ (\beta_+ - s)^{\alpha-j} + (-1)^j \gamma_- (\beta_- + s)^{\alpha-j} \right] \\ &= \Gamma(j-\alpha) \left[\gamma_+ (\beta_+)^{\alpha-j} + (-1)^j \gamma_- (\beta_-)^{\alpha-j} \right], \quad j = 2, 3, \dots \end{aligned}$$

When $\alpha - j < 0$, the terms $(\beta_+ - s)^{\alpha-j}$ and $(\beta_- + s)^{\alpha-j}$ are finite at $s = 0$ only when $\beta_+, \beta_- \neq 0$. Therefore, moments of order higher than α are finite only when both dampening coefficients (β_+, β_-) are strictly positive. When either one is zero, cumulants are finite only up to order α .

Appendix B. The Special Case of $\alpha = 1$

When $\alpha = 1$, $\Gamma(m-\alpha) = (m-2)!$. The series in (A2) and (A3) converge to different representations,

$$I_1^+ = \gamma_+ \beta_+ \sum_{m=2}^{\infty} \frac{(s/\beta_+)^m}{m(m-1)} = \gamma_+ [s + (\beta_+ - s) \ln(1 - s/\beta_+)], \quad (\text{B5})$$

$$I_1^- = \gamma_- \beta_- \sum_{m=2}^{\infty} \frac{(-s/\beta_-)^m}{m(m-1)} = \gamma_- [-s + (\beta_- + s) \ln(1 + s/\beta_-)]. \quad (\text{B6})$$

Hence,

$$\begin{aligned} k_+(s) &= \gamma_+ (\beta_+ - s) \ln(1 - s/\beta_+) + s \gamma_+ (1 + \Gamma(0, \beta_+)), \\ k_-(s) &= \gamma_- (\beta_- + s) \ln(1 + s/\beta_-) - s \gamma_- (1 + \Gamma(0, \beta_-)), \end{aligned}$$

and

$$k(s) = \gamma_+ (\beta_+ - s) \ln(1 - s/\beta_+) + \gamma_- (\beta_- + s) \ln(1 + s/\beta_-) + sC,$$

with

$$C = \gamma_+ (1 + \Gamma(0, \beta_+)) - \gamma_- (1 + \Gamma(0, \beta_-)).$$

Given the cumulant exponent, I can again derive the cumulants by progressively evaluating the derivative $\kappa_j =$

$$\left. \frac{\partial k(s)}{\partial s} \right|_{s=0},$$

$$\kappa_1 = -\gamma_+ + \gamma_- + C,$$

$$\kappa_j = (j-2)! \gamma_+ \beta_+^{-(j-1)} + (-1)^j (j-2)! \gamma_- \beta_-^{-(j-1)}, \quad j = 2, 3, \dots$$

which are finite for all j as long as $\beta_{\pm} > 0$.

References

- Aït-Sahalia, Yacine, and Andrew Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.
- Backus, David, Silverio Foresi, and Liuren Wu, 1997, Accounting for biases in Black-Scholes, Working paper, New York University.
- Barnea, Amir, and David H. Downes, 1973, A reexamination of the empirical distribution of stock price changes, *Journal of the American Statistical Association* 68, 348–350.
- Bates, David, 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181–238.
- Bates, David S., 2001, The market for crash risk, Working paper, University of Iowa IA.
- Bertoin, Jean, 1996, *Lévy Processes*. (Cambridge University Press Cambridge).
- Black, Fisher, 1976, Studies of stock price volatility changes, in *Proceedings of the 1976 American Statistical Association, Business and Economical Statistics Section* (American Statistical Association, Alexandria, VA).
- Black, Fisher, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.
- Boyarchenko, S., and S. Levendorskii, 2000, Option pricing for truncated Lévy processes, *International Journal of Theoretical and Applied Finance* 3, 549–552.
- Brenner, Menachem, 1974, On the stability of distribution of the market component in stock price changes, *Journal of Financial and Quantitative Analysis* 9, 945–961.
- Calvet, Laurent, and Adlai Fisher, 2002, Multifractality in asset returns: Theory and evidence, *Review of Economics and Statistics* 84, 381–406.
- Carr, Peter, Hélyette Geman, Dilip Madan, and Marc Yor, 2002, The fine structure of asset returns: An empirical investigation, *Journal of Business* 75, 305–332.

- Carr, Peter, Hélyette Geman, Dilip Madan, and Marc Yor, 2003, Stochastic volatility for Lévy processes, *Mathematical Finance* 13, 345–382.
- Carr, Peter, and Dilip Madan, 1999, Option valuation using the fast Fourier transform, *Journal of Computational Finance* 2, 61–73.
- Carr, Peter, and Liuren Wu, 2003, Finite moment log stable process and option pricing, *Journal of Finance* 58, 753–777.
- Carr, Peter, and Liuren Wu, 2004, Time-changed Lévy processes and option pricing, *Journal of Financial Economics* 71, 113–141.
- Chan, Terence, 1999, Pricing contingent claims on stocks driven by Lévy processes, *The Annals of Probability* 9, 504–528.
- Eberlein, Ernst, and Ulrich Keller, 1995, Hyperbolic distributions in finance, *Bernoulli* 1, 281–299.
- Engle, Robert F., and Joshua V. Rosenberg, 2002, Empirical pricing kernels, *Journal of Financial Economics* 64, 341–372.
- Fama, Eugene F., 1965, The behavior of stock market prices, *Journal of Business* 38, 34–105.
- Fielitz, Bruce D., and James P. Rozelle, 1983, Stable distributions and the mixtures of distributions hypotheses for common stock returns, *Journal of the American Statistical Association* 78, 28–36.
- Foresi, Silverio, and Liuren Wu, 2003, Crash-o-phobia: A domestic fear or a worldwide concern?, Working paper, Baruch College.
- Gerber, Hans U., and Elias S. W. Shiu, 1994, Option pricing by Esscher transforms, *Transactions of the Society of Actuaries* XLVI, 99–191.
- Haggeman, Robert, 1978, More evidence on the distribution of security returns, *Journal of Finance* 33, 1213–1221.

- Hall, Joyce A., B. Wade Brorsen, and Scott H. Irwin, 1989, The distribution of futures prices: A test of the stable paretian and mixture of normals hypotheses, *Journal of Financial and Quantitative Analysis* 24, 105–116.
- Hsu, Der-Ahn, Robert B. Miller, and Dean W. Wichern, 1974, On the stable paretian behavior of stock market prices, *Journal of the American Statistical Association* 69, 108–113.
- Huang, Jingzhi, and Liuren Wu, 2004, Specification analysis of option pricing models based on time-changed Lévy processes, *Journal of Finance* 59, 1405–1439.
- Jackwerth, Jens Carsten, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433–451.
- Jacod, Jean, and Albert N. Shiryaev, 1987, *Limit Theorems for Stochastic Processes*. (Springer-Verlag Berlin).
- Janicki, Aleksander, and Aleksander Weron, 1994, *Simulation and Chaotic Behavior of α -Stable Stochastic Processes*. (Marcel Dekker New York).
- Kallsen, Jan, 2000, Optimal portfolios for exponential Lévy processes, *Mathematical Methods of Operations Research* 51, 351–374.
- Kallsen, Jan, and Albert N. Shiryaev, 2002, The cumulant process and Esscher's change of measure, *Finance and Stochastics* 6, 397–428.
- Kou, Steven G., 2002, A jump-diffusion model for option pricing, *Management Science* 48, 1086–1101.
- Küchler, Uwe, and Michael Sørensen, 1997, *Exponential Families of Stochastic Processes*. (Springer New York).
- Madan, Dilip, and Frank Milne, 1991, Option pricing with VG martingale components, *Mathematical Finance* 1, 39–56.
- Madan, Dilip, and Eugene Seneta, 1990, The variance gamma (V.G.) model for share market returns, *Journal of Business* 63, 511–524.
- Madan, Dilip B., Peter P. Carr, and Eric C. Chang, 1998, The variance gamma process and option pricing, *European Finance Review* 2, 79–105.

- Mandelbrot, Benoit B., 1963, The variation of certain speculative prices, *Journal of Business* 36, 394–419.
- Mantegna, R. N., and H. E. Stanley, 1995, Ultra-slow convergence to a gaussian: The truncated Lévy flights, in Michael F. Shlesinger, George M. Zaslavsky, and Uriel Frisch, eds.: *Lévy Flights and Related Topics in Physics* (Springer-Verlag, Berlin).
- Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.
- Miyahara, Yoshio, 1999, Minimal entropy martingale measures of jump type price processes in incomplete assets markets, *Asia-Pacific Financial Markets* 6, 67–113.
- Officier, R. R., 1972, The distribution of stock returns, *Journal of the American Statistical Association* 67, 807–812.
- Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.
- Samorodnitsky, Gennady, and Murad S. Taqqu, 1994, *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. (Chapman & Hall New York).
- Teichmoeller, John, 1971, A note on the distribution of stock price changes, *Journal of the American Statistical Association* 66, 282–284.
- Zolotarev, Vladimir M., 1986, *One-Dimensional Stable Distributions*. (American Mathematical Society Providence, Rhode Island).

TABLE 1 Parameter Estimates of the DPL Model

Entries report the estimates, standard errors, and p -values of the model parameters. Panel A presents the estimates for the unconstrained model; panel B for the model with the constraint: $\beta_-^Q = 0$. The models are calibrated to both the time series of daily return data on S&P 500 index, from July 3, 1962 to December 31, 2001 (9942 observations), and the large cross section of S&P 500 index option prices from April 1999 to May 2000 (290 business days, 35038 observations). The calibration is based on maximum likelihood method. The last row reports the aggregate log likelihood values for the two models.

Model Parameters	<i>A. Unconstrained</i>			<i>B. Constrained</i>		
	Estimates	Std. Error	p -Value	Estimates	Std. Error	p -Value
μ	0.0776	0.0001	0.0000	0.0777	0.0001	0.0000
α	1.4892	0.0000	0.0000	1.4942	0.0000	0.0000
γ_+	0.0024	0.0000	0.0000	0.0024	0.0000	0.0000
γ_-	0.0315	0.0000	0.0000	0.0311	0.0000	0.0000
β_+	1.0015	0.0000	0.0000	1.0015	0.0000	0.0000
β_-	12.9788	0.0000	0.0000	12.9436	0.0000	0.0000
β_+^Q	5.2306	0.6282	0.0000	5.2479	0.4684	0.0000
β_-^Q	0.0067	0.0226	0.7656	0	—	—
<hr/>						
$\mathcal{L} (\times 10^4)$	-6.77767			-6.77772		