Pricing Variance Swaps on Time-Changed Lévy Processes

Peter Carr* and Roger Lee†

This version:‡ April 6, 2009

Abstract

We prove that a multiple of a log contract prices a variance swap, under arbitrary exponential Lévy dynamics, stochastically time-changed by an arbitrary continuous clock having arbitrary correlation with the driving Lévy process, subject to integrability conditions. We solve for the multiplier, which depends only on the Lévy process, not on the clock. In the case of an arbitrary continuous underlying returns process, the multiplier is 2, which recovers the standard no-jump variance swap pricing formula as a special case of our framework. In the presence of negatively-skewed jump risk, however, we prove that the multiplier exceeds 2, which agrees with calibrations of time-changed Lévy processes to equity options data. Finally we show that discrete sampling increases variance swap values, under an independence condition; so if the commonly-quoted 2 multiple undervalues the continuously-sampled variance, then it undervalues furthermore the discretely-sampled variance.

1 Introduction

A variance swap (VS) contract on an underlying price process pays (to the long party) at expiry a floating leg equal to the realized variance over the swap's fixed life, where realized variance with continuous sampling is defined as the quadratic variation of the underlying log price, and realized variance with discrete sampling is defined as the sum of squared increments of the underlying log price, typically at daily intervals. In exchange the long party pays at expiry a fixed leg, set such that the VS has zero cost of entry. Hence, a VS amounts to a forward contract on realized variance.

VS trade over-the-counter on stock indices; they also trade on single stocks (with capped payouts), and to a much lesser extent, on exchange rates and commodity futures. Highly liquid, VS on stock indices now have bid-offer spreads narrower than those of at-the-money options. VS have become the standard instrument for taking views on future realized volatility and managing volatility exposure.

*Bloomberg LP, and Courant Institute. pcarr@nyc.rr.com

†University of Chicago, Department of Mathematics. RL@math.uchicago.edu

‡We thank Tom Bielecki for helpful comments.

1.1 The ND approach

Options were first listed in the United States in 1973, just as the Black-Merton-Scholes (BMS) breakthrough for valuing options first appeared in print. For VS, the corresponding breakthrough, which we designate as the ND theory, arose in the early 1990's, first in a working paper by Neuberger [13], and then independently in a published article by Dupire [8]. VS began trading sporadically shortly thereafter, and achieved prominence in the late 1990's.

Compared to earlier efforts, the BMS option pricing formula has the advantage that it does not depend on the expected rate of return of the underlying asset. Analogously, the ND approach for VS pricing has the advantage that the ND formula does not depend on the level and dynamics of the instantaneous variance rate. The BMS formula values a vanilla option *relative* to the underlying asset (whose price incorporates the relevant information about expected returns); analogously, the ND approach values a continuously-sampled VS *relative* to a co-terminal log contract (whose price incorporates the relevant information about variance dynamics), where a log contract written on an underlying F is defined to pay $-\log(F_T/F_0)$ at its fixed maturity T.

Specifically, the ND ("Neuberger/Dupire" or "No Discontinuity") theory shows that, for a continuously-sampled VS on any underlying price process with *continuous* paths, the fair fixed payment is simply twice the forward price of the log contract. Applying an insight from Breeden and Litzenberger [2], Dupire [8] first indicated that this forward price can be obtained from co-terminal option prices at all strikes, and Carr-Madan [6] published the first explicit formula.

In 2003, the CBOE adopted a discrete implementation of this formula to revise its construction of the VIX (volatility index), a widely-quoted indicator of the options-implied expectation of short-term S&P500 realized volatility. With justification resting entirely on the ND theory, the VIX is constructed as an estimate of twice the forward price of a 30-day log contract, quoted as an annualized volatility. For the decade preceding its 2003 revision, VIX had been obtained instead from an estimate of at-the money BMS implied volatility, reflecting the prominence of the BMS model during this period. The 2003 switch, to a VS synthesized using ND theory, gave tacit recognition to the rising significance of the VS market and of the ND approach to VS pricing.

The justly celebrated ND theory, however, makes a no-jump assumption, which is restrictive especially in light of recent market events. The ND formula for VS valuation, somewhat model-free in that it holds for all continuous underlying price processes, is not completely model-free, as it can misprice VS on underlying processes which jump.

The original ND theory's applicability can be questioned in regard to its no-jump assumption, or in regard to its implications.

In regard to the former, the sharp moves experienced recently by all asset classes suggest that the ND no-jump assumption does not apply in today's markets. Moreover, even prior to the events of 2008, empirical studies concluded that option pricing models which permit jumps outperform those that assume no jumps.

In regard to the latter, the implications of ND theory can be tested in markets where one has both liquid VS quotes and accurate estimates of log contract prices. In such instances, for example the Eurostoxx or the S&P500, one can observe whether VS are quoted at twice the estimated forward price of log contracts. Although historical data from the over-the-counter VS market are scant at present, anecdotal evidence confirms that market participants indeed often observe discrepancies between market VS quotes and the ND value, and sometimes make ad-hoc adjustments in an attempt to reconcile the discrepancies unexplained by the ND theory.

1.2 Parametric approaches

Notwithstanding the widespread adoption of the nonparametric ND approach, an alternative line of research prices VS using parametric models for the underlying dynamics, typically allowing for stochastic volatility and/or jumps. For example, under CGMY dynamics for the underlying log returns, Carr-Geman-Madan-Yor [5] find pricing formulas for VS and other volatility derivatives, in terms of the CGMY model's parameters. Under Black-Scholes, Heston, Merton, and Bates dynamics, Broadie-Jain [3] find pricing formulas in terms of the respective models' parameters.

Continuous parametric models inherit the drawbacks of ND theory: a disputable assumption of no jumps, and a disputable conclusion that values continuously-sampled VS at exactly two times a log contract. Moreover, calibrations of models having finitely many parameters may be unable to achieve consistency with a full set of option price observations.

Parametric jump models do have the ability to reconcile the discrepancy between log contract and VS prices, and to price the (typically small, by [3]) effect of discrete sampling. However, they are subject to model risk. Misspecification or miscalibration of, for instance, a jump arrival rate process will generally result in erroneous VS pricing. Averse to this model risk, market participants have resisted parametric approaches to pricing VS.

1.3 Our approach

By introducing jumps in the underlying asset price, we generalize the ND theory of VS pricing. Indeed we value VS on a general exponential Lévy process, stochastically time-changed by an arbitrary unspecified continuous integrable clock. The driving Lévy process X can have jumps of finite or infinite activity, while the clock and X can have mutual dependence and correlation. Our framework includes the ND pricing theory's full scope (all positive continuous underlying prices) as a special case in which X is Brownian motion with drift. In our more general setting of time-changed Lévy processes (TCLP), we prove that a multiple of the log contract still prices the VS. We prove however that the correct multiplier is not 2 but rather a constant that depends on the characteristics of X – and *only* X. The multiplier is invariant to the time change.

Our approach makes the following contributions.

First is realism. We introduce empirically-relevant jumps into the nonparametric ND theory. Simultaneously, we introduce empirically-relevant stochastic clocks into Lévy processes, including the CGMY and Merton processes analyzed in the parametric VS literature. Stochastic clocks can generate empirical features of stock returns, such as stochastic volatility, stochastic jump arrival rates, and volatility clustering – features missing in pure Lévy models. Moreover, we allow leverage effects to arise from dependence between the clock and the Lévy driver, or from skewed jump distributions. The resulting processes are capable of achieving consistency with observed option skews at both long and short horizons.

Second is robustness. We extend, to a setting with jumps, the robustness of the ND approach to pricing VS. By declining to specify and estimate the dynamics of the clock that generates stochastic variance and jump arrival rates, we decline to price VS in terms of a full set of estimated parameters. Instead we price VS in terms of observable European option prices, using relationships valid irrespective of the time-change. We thereby avoid the model risk of misspecifying or miscalibrating the unobservable instantaneous variance and jump intensity processes.

Third is the capability to reconcile the prices of VS and log contracts. In markets where a discrepancy exists between an observable VS quote and two times the log contract valuation, the ND theory provides no mechanism to explain the observed disparity. In contrast, via choice of the driving Lévy process, our TCLP framework can achieve consistency with observations of both VS and log contracts.

Fourth is the capability to quantify the bias of ND-style VS valuations, and to explain the sign of that bias in terms of jump skewness. Using empirically calibrated TCLPs, we compute multipliers in the presence of jump risk, and find that they typically exceed the ND multiplier 2. In this setting, the VIX (modulo strike-discreteness effects) and other ND-style VS valuations therefore underestimate the risk-neutral expectation of continuously-sampled realized variance. Relating this bias to jump skewness, we show that Lévy processes has a multiplier exceeding two if and only if its Lévy measure has negative skewness, in a sense that we will define.

Fifth is the extension of nonparametric pricing to discretely-sampled VS. In practice VS contracts specify discrete (typically daily) sampling, but the VS pricing literature mainly addresses continuously-sampled variance. An exception is Broadie-Jain [3] who compute discretely-sampled VS values, in terms of parameters of particular models. In contrast, we maintain the nonparametric approach, deriving bounds for discretely-sampled VS prices in terms of log contract prices, instead of model parameters; as a corollary, we show that discrete sampling increases the value of variance swaps, under an independence condition. Hence if ND theory undervalues the continuously-sampled VS, then in this setting it undervalues furthermore the discretely-sampled VS.

The remainder of this paper is organized as follows. Section 2 introduces and characterizes the *multiplier* of a Lévy process. Section 3 proves that the fair fixed payment on the VS is just the multiplier times the forward price of the log contract, where the multiplier depends only on the driving Lévy process, not on the time change. Section 4 calculates multiplier formulas for several popular Lévy processes, it presents evidence that empirically calibrated TCLP's produce multipliers which exceed two (hence VS prices which exceed ND valuations), and it relates this phenomenon to negative skewness. Section 5 analyzes the impact of discrete sampling. Section 6 concludes.

2 The Multiplier

We work in a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_u\}_{u\geq 0}, \mathbb{P})$.

Let brackets $[\cdot]$ denote quadratic variation.

Proposition 2.1. *Let L be a Lévy process with Lévy measure ν and Brownian variance σ^2. Then*

$$\mathbb{E}[L]_1 = \sigma^2 + \int x^2 \mathrm{d}\nu(x) \in [0, \infty].$$

Moreover, $\mathbb{E}[L]_1 < \infty$ if and only if $\mathbb{E}L_1^2 < \infty$.

Proof. We have

$$[L]_1 = \sigma^2 + \sum_{0<u\leq 1} (\Delta L_u)^2.$$

Sato [16] Propositions 19.2 and 19.5, applied to the restriction of ν to $\{x : 1/m < |x| < m\}$ for each $m > 0$, together with monotone convergence as $m \to \infty$, imply that

$$\mathbb{E}[L]_1 = \sigma^2 + \int x^2 \mathrm{d}\nu(x).$$

So the following conditions are equivalent:

$$\mathbb{E}[L]_1 < \infty \quad \Leftrightarrow \quad \int x^2 \mathrm{d}\nu(x) < \infty \quad \Leftrightarrow \quad \mathbb{E}L_1^2 < \infty,$$

where the last step is by Sato [16] Corollary 25.8 and the general property $\int_{|x|<1} x^2 \mathrm{d}\nu(x) < \infty$. $\square$

The following corollary is immediate.

Corollary 2.2. *If $\mathbb{E}[L]_1 < \infty$ then $\mathbb{E}|L_1| < \infty$.*

Let us define the *multiplier* of a returns-driving process.

Definition 2.3 (Returns-driving process)**.** *A returns-driving process is a nonconstant Lévy process X such that $\mathbb{E}e^{X_1} < \infty$ and $\mathbb{E}[X]_1 < \infty$.*

Definition 2.4 (Multiplier)**.** *Define the* multiplier *of a returns-driving process X by*

$$Q_X := \frac{\mathbb{E}[X]_1}{\log \mathbb{E}e^{X_1} - \mathbb{E}X_1}.$$

Proposition 2.5. *For any returns-driving process X, the multiplier exists and satisfies*

$$0 < Q_X = \frac{\operatorname{Var} X_1}{\log \mathbb{E}e^{X_1} - \mathbb{E}X_1} = \frac{\kappa_X''(0)}{\kappa_X(1) - \kappa_X'(0)}.$$

where $\kappa_X(z) := \log \mathbb{E}e^{zX_1}$ denotes the cumulant generating function of X, and primes denote right derivatives.

Proof. The multiplier is well-defined and positive because $\mathbb{E}e^{X_1} > e^{\mathbb{E}X_1}$ by convexity of exp and nonconstancy of X.

Define the martingale $M_u := X_u - u\mathbb{E}X_1$. The middle Q_X formula follows from rewriting the numerator as

$$\mathbb{E}[X]_1 = \mathbb{E}[M]_1 = \mathbb{E}M_1^2 = \operatorname{Var} X_1,$$

where the middle inequality is by, for instance, Protter [14] Corollary 27.3.

The final Q_X formula follows from the existence of κ_X on $[0,1]$. □

Proposition 2.6. *Let X be a returns-driving process with generating triplet (σ^2, ν, γ). Then*

$$Q_X = \frac{\sigma^2 + \int x^2 \nu(\mathrm{d}x)}{\sigma^2/2 + \int (e^x - 1 - x)\nu(\mathrm{d}x)}.$$

Proof. Sato [16] Theorem 25.17 implies that

$$\log \mathbb{E}e^{X_1} = \sigma^2/2 + \int (e^x - 1 - x\mathbb{I}_{|x|\leq 1})\nu(\mathrm{d}x) + \gamma$$

(and that the integral is finite). Sato [16] Example 25.12 implies that

$$-\mathbb{E}X_1 = -\gamma - \int_{|x|\geq 1} x\nu(\mathrm{d}x)$$

(and that the integral is finite). Summing gives the denominator of Q_X.

Proposition 2.1 gives the numerator. □

3 Variance Swaps and Log Contracts

This section's assumptions will apply throughout the remainder of this paper.

Fix a time horizon $T > 0$.

Let the interest rate be a deterministic right-continuous process r such that $\int_0^T |r_s|\mathrm{d}s < \infty$. Let

$$R_t := \int_0^t r_s \mathrm{d}s.$$

Let F denote a positive underlying T-expiry forward or futures price process, and let

$$Y_t := \log(F_t/F_0)$$

denote the log return on F. Let

$$F_t^* := F_t e^{R_t - R_T}.$$

denote the corresponding underlying spot price, and

$$Y_t^* := \log(F_t^*/F_0^*) = Y_t + R_t$$

denote the log return on F^*.

Define the T-expiry *log contract* to pay at time T

$$-Y_T,$$

where the sign convention conveniently makes log contracts have nonnegative value. Define the (floating leg of a continuously-sampled) *variance swap* on F pay at time T

$$[Y]_T.$$

Assume that

$$Y_t = \bar{X}_{\tau_t} \tag{3.1}$$

where

$$\bar{X}_u := X_u - u \log \mathbb{E}e^{X_1} \tag{3.2}$$

for some returns-driving process X in the sense of Definition 2.3, and where the time change

$$\{\tau_t : t \in [0, T]\}$$

is a continuous increasing family of stopping times. We do *not* assume independence of X and τ.

Financially, we regard X, indexed by "business" time, as a "driving" or "background" Lévy process, which induces the drift-adjusted process $\bar{X}$ such that $e^{\bar{X}}$ is a martingale. We regard τ as an unspecified stochastic clock that maps calendar time t to business time τ_t. The resulting $\{\mathcal{F}_{\tau_t}\}$-adapted process Y can exhibit stochastic volatility, stochastic jump-intensity, volatility clustering, and "leverage" effects, the latter via skewed jump distributions, or via correlation of X and τ.

Assume that $\mathbb{P}$ is a martingale measure for log contracts and variance swaps; in particular assume that the T-expiry log contract and continuously-sampled variance swap have respective time-0 values $e^{-R_T}\mathbb{E}(-Y_T)$ and $e^{-R_T}\mathbb{E}[Y]_T$, if finite.

Proposition 3.1 (Variance swap valuation)**.** *If* $\mathbb{E}\tau_T < \infty$ *then*

$$\mathbb{E}[Y]_T = Q_X \mathbb{E}(-Y_T).$$

The multiplier Q_X *does not depend on the time-change.*

Proof. The definition of Q_X and the equality $[X] = [\bar{X}]$ imply that the Lévy process $[\bar{X}]_u + Q_X \bar{X}_u$ is a martingale. Because $\mathbb{E}\tau_T < \infty$, we have, by Wald's first equation in continuous time [9],

$$\mathbb{E}\big([\bar{X}]_{\tau_T} + Q_X \bar{X}_{\tau_T}\big) = 0.$$

Moreover $\mathbb{E}[\bar{X}]_{\tau_T} < \infty$, again by Wald's first equation, so

$$\mathbb{E}[\bar{X}]_{\tau_T} = Q_X \mathbb{E}(-\bar{X}_{\tau_T}).$$

Finally, the continuity of τ implies $[Y]_T = [\bar{X}]_{\tau_T}$, by Jacod [10] Theorem 10.17. □

Hence the variance swap value $e^{-R_T}\mathbb{E}[Y]_T$ equals Q_X times the log contract value $e^{-R_T}\mathbb{E}(-Y_T)$. Equivalently, restated in terms of forward-settled payments, the variance swap fixed payment's fair level $\mathbb{E}[Y]_T$ equals Q_X times the log contract's forward price $\mathbb{E}(-Y_T)$.

The *multiplier* Q_X depends only on the characteristics of the driving Lévy process. It does *not* depend on the time-change.

Likewise, for the spot underlying, the (floating leg of a continuously-sampled) variance swap on F^* can be defined to pay $[Y^*]_T$. However, $[Y] = [Y^*]$ because $Y^* - Y = R$ has finite variation and no jumps. Therefore, no distinction exists between (continuously-sampled) variance swaps on futures and spot. We have established the following.

Corollary 3.2 (Variance swap valuation, on spot underlying)**.** *If* $\mathbb{E}\tau_T < \infty$ *then*

$$\mathbb{E}[Y^*]_T = Q_X \mathbb{E}(-Y_T).$$

4 Multiplier Calculations

In the following examples of returns-driving processes X, we will not need to specify the "drift" component of X, because passing to $\bar{X}$ via (3.2) resets the drift anyway, to make $e^{\bar{X}}$ a martingale.

We emphasize that each example's scope includes a *family* of log returns processes $Y_t = \bar{X}_{\tau_t}$, because the time change τ is general and unspecified. Without modeling the stochastic clock τ, Proposition 3.1 prices the variance swap payoff $[Y]_T$ in each case.

4.1 Example: Time-changed Brownian motion

Let X be Brownian motion. Then

$$Q_X = \frac{\mathbb{E}[X]_1}{\log \mathbb{E}e^{X_1} - \mathbb{E}X_1} = \frac{1}{1/2} = 2. \tag{4.1}$$

This multiplier prices variance swaps on all positive continuous local martingales, because their log return dynamics are all generated by time changes of drift-adjusted Brownian motion:

Proposition 4.1. *Let S be a positive continuous local martingale relative to a filtration $(\mathcal{G}_t)_{t\geq 0}$. If $\mathbb{E}[\log S]_T < \infty$ and $[\log S]_\infty = \infty$, then there exist a filtration $(\mathcal{F}_u)_{u\geq 0}$, an $\mathcal{F}$-Brownian motion W, and a continuous $\mathcal{F}$-time change τ with $\mathbb{E}\tau_T < \infty$, such that $\log(S_t/S_0) = W_{\tau_t} - \tau_t/2$.*

Proof. We have

$$\mathrm{d}\log S_t = \frac{1}{S_t}\mathrm{d}S_t - \frac{1}{2S_t^2}\mathrm{d}[S]_t = \frac{1}{S_t}\mathrm{d}S_t - \frac{1}{2}\mathrm{d}[\log S]_t \tag{4.2}$$

hence

$$M_t := \int_0^t \frac{1}{S_t}\mathrm{d}S_t = \log(S_t/S_0) + \frac{1}{2}[\log(S_./S_0)]_t \tag{4.3}$$

is a continuous local martingale.

Define $A_u := \inf\{t : [M]_t \geq u\}$ and $\mathcal{F}_u := \mathcal{G}_{A_u}$ and $\tau_t := [M]_t = [\log(S_./S_0)]_t$.

By Dambis/Dubins-Schwarz, $W_u := M_{A_u}$ is $\mathcal{F}$-Brownian motion, and τ is an $\mathcal{F}$-time change, and $W_{\tau_t} = M_t$. Hence $\log(S_t/S_0) = W_{\tau_t} - \tau_t/2$ as claimed. □

The assumption that $[\log S]_\infty = \infty$ can be removed by enlargement of the probability space; see for example Revuz-Yor [15] Theorem V.1.7.

Consequently, our Proposition 3.1 includes as a special case the classical price equivalence of a variance swap and 2 log contracts, for all continuous underlying log returns processes, because all such dynamics arise via (3.1)-(3.2) from some Brownian X, according to Proposition 4.1.

Proposition 3.1 extends the classical result by allowing general time-changes of *general* Lévy processes X.

4.2 Example: Time-changed jump diffusion, with fixed jump sizes

Let X have Brownian variance σ^2 and Lévy measure

$$\lambda_1\delta_{c_1} + \lambda_2\delta_{c_2}, \tag{4.4}$$

where δ_c denotes a point mass at c. Let $c_1 > 0$ and $c_2 < 0$; thus up-jumps have magnitude c_1 and down-jumps have magnitude $|c_2|$.

By Proposition 2.6, the multiplier is exactly

$$Q_X = \frac{\sigma^2 + \lambda_1 c_1^2 + \lambda_2 c_2^2}{\sigma^2/2 + \lambda_1(e^{c_1} - 1 - c_1) + \lambda_2(e^{c_2} - 1 - c_2)}. \tag{4.5}$$

If $\sigma \neq 0$, then a third-order Taylor approximation in (c_1, c_2) about $(0,0)$ gives the approximation

$$Q_X \approx 2 - \frac{2\lambda_1}{3}c_1^3 + \frac{2\lambda_2}{3}|c_2|^3. \tag{4.6}$$

which is increasing in absolute down-jump size, but decreasing in up-jump size.

4.3 Example: Time-changed Kou double-exponential jump diffusion

Let X be a Kou [11] process, defined by Brownian variance σ^2 and Lévy density

$$\nu(x) = \lambda_1 a_1 e^{-a_1|x|}\mathbb{I}_{x>0} + \lambda_2 a_2 e^{-a_2|x|}\mathbb{I}_{x<0} \tag{4.7}$$

where $a_1 \geq 1$ and $a_2 > 0$. Hence up-jumps and down-jumps have mean absolute size $1/a_1$ and $1/a_2$ respectively.

By Proposition 2.6, the multiplier is exactly

$$Q_X = \frac{\sigma^2 + 2\lambda_1/a_1^2 + 2\lambda_2/a_2^2}{\sigma^2/2 + \lambda_1/(a_1-1) - \lambda_2/(a_2+1) - \lambda_1/a_1 + \lambda_2/a_2}. \tag{4.8}$$

If $\sigma \neq 0$, then a third-order Taylor expansion in $(1/a_1, 1/a_2)$ about $(0,0)$ gives the approximation

$$Q_X \approx 2 - \frac{4\lambda_1/\sigma^2}{a_1^3} + \frac{4\lambda_2/\sigma^2}{a_2^3}, \tag{4.9}$$

which is increasing in mean absolute down-jump size, but decreasing in mean up-jump size.

The case that price trajectories F_t are piecewise constant (changing only at jump times) corresponds to $(\sigma^2, \lambda_1, \lambda_2) = (0, a_1 - 1, a_2 + 1)$, times any positive scalar. In this pure-jump case, the multiplier (4.8) becomes exactly

$$Q_X = 2 - \frac{2}{a_1} + \frac{2}{a_2}, \tag{4.10}$$

which is 2, minus twice the mean up-jump size, plus twice the mean absolute down-jump size.

4.4 Example: Time-changed Merton lognormal jump diffusion

Let X be a Merton [12] jump-diffusion, defined by Brownian variance σ^2 and Lévy density

$$\nu(x) = \frac{\lambda}{\eta\sqrt{2\pi}} \exp\Big(\frac{-(x-\mu)^2}{2\eta^2}\Big). \tag{4.11}$$

By Proposition 2.6, the multiplier is exactly

$$Q_X = \frac{\sigma^2 + \lambda\eta^2 + \lambda\mu^2}{\sigma^2/2 + \lambda(e^{\mu+\eta^2/2} - 1 - \mu)}. \tag{4.12}$$

If $\sigma \neq 0$, then a third-order Taylor expansion in (μ, η) about $(0, 0)$ gives the approximation

$$Q_X \approx 2 - \frac{2\lambda}{\sigma^2}\eta^2\mu - \frac{2\lambda}{3\sigma^2}\mu^3, \tag{4.13}$$

which is decreasing in μ.

4.5 Example: Time-changed CGMY

Let X have no Brownian component. Let X have the *generalized CGMY* Lévy density

$$\nu(x) = \frac{C_n}{|x|^{1+Y_n}} e^{-G|x|}\mathbb{I}_{x<0} + \frac{C_p}{|x|^{1+Y_p}} e^{-M|x|}\mathbb{I}_{x>0}, \tag{4.14}$$

where $C_p, C_n > 0$, and $G, M > 0$, and $Y_p, Y_n < 2$.

Then, in the case $\{Y_p, Y_n\} \cap \{0, 1\} = \emptyset$, the cumulant generating function of X is

$$\log \mathbb{E}e^{zX_1} = \gamma z + C_p\Gamma(-Y_p)[(M-z)^{Y_p} - M^{Y_p}] + C_n\Gamma(-Y_n)[(G+z)^{Y_n} - G^{Y_n}], \tag{4.15}$$

for some γ that we need not specify. By Proposition 2.5, the multiplier for CGMY is exactly

$$Q_X = \frac{C_n\Gamma(2-Y_n)G^{Y_n-2} + C_p\Gamma(2-Y_p)M^{Y_p-2}}{C_n\Gamma(-Y_n)[(G+1)^{Y_n} - G^{Y_n} - Y_nG^{Y_n-1}] + C_p\Gamma(-Y_p)[(M-1)^{Y_p} - M^{Y_p} + Y_pM^{Y_p-1}]}. \tag{4.16}$$

Expanding the denominator in $1/G$ and $1/M$ yields

$$Q_X \approx 2 \times \frac{C_n\Gamma(2-Y_n)G^{Y_n-2} + C_p\Gamma(2-Y_p)M^{Y_p-2}}{C_n\Gamma(2-Y_n)G^{Y_n-2}(1 - \frac{2-Y_n}{3G} + \ldots) + C_p\Gamma(2-Y_p)M^{Y_p-2}(1 + \frac{2-Y_p}{3M} + \ldots)}. \tag{4.17}$$

The basic CGMY model takes $C_p = C_n$ and $Y_p = Y_n = Y$. Its multiplier is exactly

$$Q_X = \frac{Y(1-Y)(G^{Y-2}+M^{Y-2})}{G^Y - (G+1)^Y + YG^{Y-1} + M^Y - (M-1)^Y - YM^{Y-1}} \tag{4.18}$$

and approximately

$$Q_X \approx 2 \times \frac{G^{Y-2}+M^{Y-2}}{G^{Y-2}(1-\frac{2-Y}{3G}+\ldots) + M^{Y-2}(1+\frac{2-Y}{3M}+\ldots)}, \tag{4.19}$$

which reveals a sign asymmetry between the $-\frac{2-Y}{3G}$ and the $+\frac{2-Y}{3M}$.

4.6 Example: Time-changed VG

Taking $Y = 0$ in (4.14) produces the *variance gamma* (VG) process, which replaces (4.15) with

$$\log \mathbb{E} e^{zX_1} = \gamma z - C\log(1+z/G) - C\log(1-z/M). \tag{4.20}$$

By Proposition 2.5, the multiplier for VG is exactly

$$Q_X = \frac{1/G^2 + 1/M^2}{1/G - \log(1+1/G) - 1/M - \log(1-1/M)}. \tag{4.21}$$

Expanding the denominator in $1/G$ and $1/M$ yields the $Y = 0$ case of the approximation (4.19).

4.7 Example: Time-changed NIG

Let X have no Brownian component. Let X have the *normal inverse Gaussian* (NIG) Lévy density

$$\nu(x) = \frac{\delta\alpha}{\pi}\frac{\exp(\beta x)K_1(\alpha|x|)}{|x|}, \tag{4.22}$$

where $\delta > 0$, $\alpha > 0$, $|\beta| < \alpha$, and K_1 denotes the modified Bessel function of the second kind and order 1. Then X has cumulant generating function

$$\log \mathbb{E} e^{zX_1} = \gamma z + \delta(\sqrt{\alpha^2-\beta^2} - \sqrt{\alpha^2-(\beta+z)^2}), \tag{4.23}$$

for some γ that we need not specify. By Proposition 2.5, the multiplier is exactly

$$Q_X = \frac{\alpha^2/(\alpha^2-\beta^2)}{\alpha^2-\beta^2-\beta-\sqrt{(\alpha^2-\beta^2)(\alpha^2-(\beta+1)^2)}}. \tag{4.24}$$

The small jump-size limit is obtained by taking $\alpha \to \infty$ which concentrates the Lévy measure near 0. Expanding in $1/\alpha$ yields the multiplier approximation

$$Q_X \approx 2 - \frac{4\beta+1}{2\alpha^2}. \tag{4.25}$$

which is decreasing in β, the parameter which controls the "tilt" of the NIG distribution.

4.8 Up-jumps, down-jumps, and skewness

The multiplier approximations (4.6) and (4.9) and (4.13) and (4.25) all exhibit a common theme: increasing the sizes of *up*-jumps (by increasing c_1 or $1/a_1$ or μ or β respectively) has the leading-order effect of *decreasing* the multiplier, whereas increasing the sizes of *down*-jumps (by increasing $|c_2|$ or $1/a_2$ or decreasing μ or β) has the leading-order effect of *increasing* the multiplier. Likewise, in the multiplier (4.19), taking larger up-jumps via $(G, M) = (B, b)$ where $B > b$, gives a smaller multiplier than taking larger down-jumps by swapping $(G, M) = (b, B)$.

Remark 4.1. This asymmetry can be explained as follows. Under any of those dynamics, we have

$$\begin{aligned} -2\log(F_T/F_0) &= \int_{0+}^{T} \frac{-2}{F_{t-}} \mathrm{d}F_t + \frac{1}{2}\int_{0+}^{T} \frac{2}{F_{t-}^2} \mathrm{d}[F]_t^c + \sum_{0<t\leq T} \Big(-2\Delta \log F_t - \frac{-2}{F_{t-}}\Delta F_t \Big) \\ &= \int_{0+}^{T} \frac{-2}{F_{t-}} \mathrm{d}F_t + [Y]_T + \sum_{0<t\leq T} \Big(\frac{2}{F_{t-}}\Delta F_t - 2\Delta Y_t - (\Delta Y_t)^2 \Big). \end{aligned} \tag{4.26}$$

where $[F]^c$ denotes the continuous part of the quadratic variation.

So 2 log contract payoffs, together with the zero-expectation profit/loss from dynamically holding $2/F_{t-}$ futures, replicate

$$[Y]_T + \sum_{0<t\leq T} \Big(2e^{\Delta Y_t} - 2 - 2\Delta Y_t - (\Delta Y_t)^2 \Big) \;\approx\; [Y]_T + \sum_{0<t\leq T} \frac{1}{3}(\Delta Y_t)^3. \tag{4.27}$$

Therefore, in the presence of up-jumps ($\Delta Y_t > 0$), the intuition is $2\mathbb{E}(-\log F_T/F_0) > \mathbb{E}[Y]_T$, and hence the 2 should be decreased in order to achieve equality, whereas in the presence of down-jumps ($\Delta Y_t < 0$), the inequality is reversed, and hence the 2 should be increased.

The calculations (4.26) and (4.27) resemble closely the jump analysis by Derman et al. [7], but the conclusion differs, because Derman et al. consider contracts which define the realized variance of a jump to be $(\Delta F_t/F_{t-})^2$ instead of $(\Delta \log F_t)^2$, which affects the leading (cubic) term.

Motivated by Remark 4.1 and Proposition 2.6, we define a relevant notion of skewness.

Definition 4.2 (Exponential skewness)**.** *For a Lévy measure ν such that $\int_{|x|>1} e^x \nu(\mathrm{d}x) < \infty$ and $\int_{|x|>1} x^2 \nu(\mathrm{d}x) < \infty$, define the* exponential skewness *of ν by*

$$6\int (e^x - 1 - x - x^2/2)\nu(\mathrm{d}x). \tag{4.28}$$

Rewriting exponential skewness as $\int (x^3 + x^4/4 + x^5/20 + \cdots)\nu(\mathrm{d}x)$ shows that the leading term of exponential skewness equals the third moment of the Lévy measure.

The connection between exponential skewness and the multiplier is proved as follows.

Proposition 4.3. *For any returns-driving process X with Lévy measure ν, we have $Q_X > 2$ if and only if ν has negative exponential skewness.*

Proof. By (4.28), exponential skewness is negative if and only if

$$\sigma^2/2 + \int (e^x - 1 - x)\nu(\mathrm{d}x) < \sigma^2/2 + \int (x^2/2)\nu(\mathrm{d}x) \tag{4.29}$$

where σ^2 denotes the Brownian variance of X. By Proposition 2.6, this is equivalent to $Q_X > 2$. $\square$

In this sense, negatively-skewed Lévy processes have multipliers greater than 2.

4.9 Multipliers of empirically estimated processes

Carr-Geman-Madan-Yor [4] calibrate various time-changed Lévy processes to data. In Table 1 we compute the multipliers associated with the parameter estimates.

In each case, the time-change is by a CIR process. We do not report the estimated parameters of the time changes, because the multiplier depends only on the driving Lévy process. The multipliers implicit in the Carr-Geman-Madan-Yor data fall in the range 2.15±0.06, except for two observations near 2.40.

Using a multiplier of 2 (or smaller) in the presence of jumps would in most cases underestimate the expectation of quadratic variation by 5 to 10 percent, and in two cases by around 20 percent.

Table 1: Carr-Geman-Madan-Yor calibration, using 4 cross-sections of S&P 500 options data.

Lévy process	Data	Lévy parameters	Multiplier
CGMY	Mar 2000	$C_n/C_p = 0.2883, G = 0.697, M = 22.0, Y_p = -3.65, Y_n = 1.45$	2.43
VG	Mar 2000	$G = 7.33, M = 32.4$	2.17
NIG	Mar 2000	$\alpha = 96.4, \beta = -92.0$	2.21
CGMY	Jun 2000	$C_n/C_p = 0.0526, G = 0.423, M = 24.6, Y_p = -4.51, Y_n = 1.67$	2.37
VG	Jun 2000	$G = 11.0, M = 30.1$	2.10
NIG	Jun 2000	$\alpha = 69.7, \beta = -62.1$	2.12
CGMY	Sep 2000	$C_n/C_p = 0.0676, G = 1.64, M = 16.9, Y_p = -2.90, Y_n = 1.54$	2.17
VG	Sep 2000	$G = 12.4, M = 33.6$	2.09
NIG	Sep 2000	$\alpha = 99.8, \beta = -91.1$	2.11
CGMY	Dec 2000	$C_n/C_p = 0.0855, G = 3.68, M = 52.9, Y_p = -2.12, Y_n = 1.22$	2.13
VG	Dec 2000	$G = 11.7, M = 42.7$	2.10
NIG	Dec 2000	$\alpha = 274.8, \beta = -265.4$	2.10

5 Discrete Sampling

Consider an arbitrary sequence of fixed sampling times $0 = t_0 < t_1 < \cdots < t_N = T$.

For $n = 0, \ldots, N-1$, and any stochastic process Z, write $\Delta_n Z := Z_{t_{n+1}} - Z_{t_n}$.

Define the payoffs of a *discretely sampled variance swap* on futures F and spot F^* to be, respectively,

$$V_T := \sum_{n=0}^{N-1} (\Delta_n Y)^2 \tag{5.1}$$

$$V_T^* := \sum_{n=0}^{N-1} (\Delta_n Y^*)^2 = \sum_{n=0}^{N-1} (\Delta_n Y + \Delta_n R)^2. \tag{5.2}$$

Unlike the continuous-sampling payoffs which satisfied $[Y] = [Y^*]_T$, the discrete-sampling payoffs V_T and V_T^* are not generally equal.

Broadie-Jain [3] express the price $\mathbb{E}V_T^*$ in terms of the parameters of some particular models; whereas we maintain a nonparametric approach which bounds $\mathbb{E}V_T^* - \mathbb{E}[Y^*]_T$ in terms of log contract prices, instead of model parameters. Bondarenko [1] values a claim which pays a non-standard definition of discretely-sampled variance (which he defends); whereas we adhere to the standard definition V_T^* of discretely-sampled variance, as referenced in actual variance swap contracts.

Still working under the Section 3 framework, let $\mathbb{E}(\cdot|\tau)$ denote expectation conditional on the σ-algebra generated by $\{\tau_t : t \le T\}$. The following formula links the discretely-sampled variance swap value $\mathbb{E}V_T^*$ back to the continuously-sampled variance swap value $\mathbb{E}[Y^*]_T$, which is already understood via Corollary 3.2. In this section (and not in any other section), our results assume independence of X and τ.

Proposition 5.1 (Variance swap valuation, with discrete sampling of spot)**.** *Assume that* $\mathbb{E}\tau_T < \infty$*, and that* τ *and* X *are independent. Then*

$$\mathbb{E}V_T^* = \mathbb{E}[Y^*]_T + \sum_{n=0}^{N-1} (\mathbb{E}\Delta_n Y^*)^2 + \mathrm{Var}(\mathbb{E}(Y_T^*|\tau)). \tag{5.3}$$

The last term has explicit form

$$\mathrm{Var}(\mathbb{E}(Y_T^*|\tau)) = (\mathbb{E}X_1)^2 \; \mathbb{E} \sum_{n=0}^{N-1} (\Delta_n \tau)^2 - \sum_{n=0}^{N-1} (\mathbb{E}\Delta_n Y)^2. \tag{5.4}$$

Proof. The Lévy process

$$L_u := X_u - \mathbb{E}X_u = X_u - u\mathbb{E}X_1 \tag{5.5}$$

is a martingale, and $\mathbb{E}\tau_T < \infty$ implies that for $t \in [0, T]$,

$$M_t := L_{\tau_t} = Y_t - \tau_t \mathbb{E}X_1 \tag{5.6}$$

is a martingale. Then, for each n, abbreviating the Δ_n notation as Δ, we have

$$\mathbb{E}(\Delta[Y^*]) = \mathbb{E}(\Delta[Y]) = \mathbb{E}(\Delta[M]) = \mathbb{E}(\Delta M)^2 = \mathbb{E}(\Delta Y - (\Delta\tau)\mathbb{E}X_1)^2 \tag{5.7}$$

where the second equality is because $Y_t - M_t = \tau_t \mathbb{E}X_1$ has finite variation and no jumps, and the third equality follows from Protter [14] Corollary 27.3. By the independence condition,

$$\mathbb{E}(\Delta Y|\tau) = (\Delta\tau)\mathbb{E}X_1, \tag{5.8}$$

which implies that (5.7) becomes

$$\begin{aligned}\mathbb{E}(\Delta[Y^*]) &= \mathbb{E}(\Delta Y - \mathbb{E}(\Delta Y|\tau))^2 \\ &= \mathbb{E}(\mathrm{Var}(\Delta Y|\tau)) \\ &= \mathrm{Var}(\Delta Y) - \mathrm{Var}(\mathbb{E}(\Delta Y|\tau)) \\ &= \mathbb{E}(\Delta Y^*)^2 - (\mathbb{E}\Delta Y^*)^2 - \mathrm{Var}(\mathbb{E}(\Delta Y|\tau))\end{aligned} \tag{5.9}$$

Sum from $n = 0$ to $N - 1$. By the τ-conditional independence of $\Delta_n Y$ and $\Delta_m Y$ for $m \neq n$,

$$\sum_{n=0}^{N-1} \mathrm{Var}(\mathbb{E}(\Delta_n Y|\tau)) = \mathrm{Var}(\mathbb{E}(Y_T|\tau)) = \mathrm{Var}(\mathbb{E}(Y_T^*|\tau)) \tag{5.10}$$

which proves (5.3). Moreover, by (5.8), we have $\mathrm{Var}(\mathbb{E}(\Delta Y|\tau)) = \mathbb{E}(\Delta\tau\mathbb{E}X_1)^2 - (\mathbb{E}\Delta Y)^2$, which proves (5.4). □

Corollary 5.2 (Variance swap lower bound, with discrete sampling of spot)**.** *Under the assumptions of Proposition 5.1, we have*

$$\mathbb{E}V_T^* \geq \mathbb{E}[Y^*]_T + \sum_{n=0}^{N-1} (\mathbb{E}Y_{t_{n+1}}^* - \mathbb{E}Y_{t_n}^*)^2. \tag{5.11}$$

The lower bound is observable *via the prices of log contracts at expiries* $t_1, \ldots, t_N$*.*

Equality holds *if the time change* τ *is nonrandom.*

Proof. This follows from (5.3) and $\mathrm{Var}(\mathbb{E}(Y_T^*|\tau)) \geq 0$ with equality if τ is nonrandom. □

Proposition 5.3 (Discrete sampling of futures)**.** *Proposition 5.1 and Corollary 5.2 apply to discretely sampled variance swaps on futures, by deleting all instances of stars* (*) *in the statements.*

Proof. The proofs still stand after deleting all instances of stars. □

Discrete sampling, therefore, increases the value of a variance swap, under an independence condition. So if the commonly-quoted 2 multiple undervalues the continuously-sampled variance swap (as suggested by the data in Section 4.9), then in this setting the 2 multiple furthermore undervalues the discretely-sampled variance swap.

6 Conclusions

Assuming continuous underlying price paths, the standard theory shows that a variance swap has the same value as two log contracts on the underlying. This valuation formula provides a standard reference point for volatility traders, and forms the basis of popular volatility indicators such as the VIX, VXN, and VDAX-NEW. However, the continuity assumption is empirically rejected in equity markets. This motivates our analysis of jump processes.

We generalize the underlying dynamics to arbitrary time-changed exponential Lévy processes (under integrability conditions), where the background Lévy process may have jumps of arbitrary distribution, and where the stochastic time-change, an arbitrary continuous clock, may have arbitrary dependence or correlation with the Lévy process. This allows stochastic volatility, stochastic jump intensity, volatility clustering, and leverage effects.

We prove that a multiple of the log contract still prices the variance swap. The *multiplier*, not necessarily 2 in this general setting, depends only on the characteristics of the driving Lévy process, not on the time change.

We calculate explicitly the multiplier for various examples of driving Lévy processes. We recover the standard no-jump valuation formula as a special case, because all positive continuous martingales are time-changes of driftless geometric Brownian motion, which has multiplier 2. We then solve for jump dynamics, including time-changes of CGMY, VG, NIG, Kou, Merton, and fixed-jump-size processes.

We observe that increasing the sizes of up-jumps tends to decrease the multiplier, whereas increasing the sizes of down-jumps tends to increase the multiplier. More precisely, we show that the multiplier exceeds 2 if and only if the jumps have negative skewness in a sense that we define. We compute, moreover, the multipliers associated with published empirical calibrations of time-changed Lévy processes, and obtain results in the range 2.1 to 2.4, which is consistent with negatively-skewed jump risk.

Finally we prove that discrete sampling increases variance swap values, under an independence condition. So if the commonly-quoted 2 multiple undervalues the continuously-sampled variance swap (as suggested by our multiplier estimates of greater than 2.1), then in this setting the 2 multiple undervalues, furthermore, the discretely-sampled variance swap.

References

[1] Oleg Bondarenko. Variance trading and market price of variance risk. University of Illinios at Chicago, 2007.

[2] David Breeden and Robert Litzenberger. Prices of state contingent claims implicit in options prices. *Journal of Business*, 51:621–651, 1978.

[3] Mark Broadie and Ashish Jain. The effect of jumps and discrete sampling on volatility and variance swaps. *International Journal of Theoretical and Applied Finance*, 11(8):761–797, 2008.

[4] Peter Carr, Hélyette Geman, Dilip Madan, and Marc Yor. Stochastic volatility for Lévy processes. *Mathematical Finance*, 13(3):345–382, 2003.

[5] Peter Carr, Hélyette Geman, Dilip Madan, and Marc Yor. Pricing options on realized variance. *Finance and Stochastics*, 9(4):453–475, 2005.

[6] Peter Carr and Dilip Madan. Towards a theory of volatility trading. In R. Jarrow, editor, *Volatility*, pages 417–427. Risk Publications, 1998.

[7] Emanuel Derman, Kresimir Demeterfi, Michael Kamal, and Joseph Zou. A guide to volatility and variance swaps. *Journal of Derivatives*, 6(4):9–32, 1999.

[8] Bruno Dupire. Model art. *Risk*, 6(9):118–124, 1993.

[9] W. J. Hall. On Wald's equations in continuous time. *Journal of Applied Probability*, 7:59–68, 1970.

[10] Jean Jacod. *Calcul Stochastique et Problèmes de Martingales*. Lecture Notes in Mathematics. Springer-Verlag, 1979.

[11] Steven Kou. A jump diffusion model for option pricing. *Management Science*, 48(8):1086–1101, 2002.

[12] Robert C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3:125–144, 1976.

[13] Anthony Neuberger. Volatility trading. London Business School working paper, 1990.

[14] Philip Protter. *Stochastic Integration and Differential Equations*. Springer-Verlag, 2nd edition, 2004.

[15] Daniel Revuz and Marc Yor. *Continuous Martingales and Brownian Motion*. Springer, 3rd edition, 1999.

[16] Ken-Iti Sato. *Lévy Processes and Infinitely Divisible Distributions*. Cambridge University Press, 1999.