Finance Stoch
DOI 10.1007/s00780-011-0157-9

Variance swaps on time-changed Lévy processes

Peter Carr · Roger Lee · Liuren Wu

Received: 10 April 2010 / Accepted: 10 February 2011

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. 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. Moreover, we show that discrete sampling increases variance swap values, under an independence condition; so if the commonly quoted multiple 2 undervalues the continuously sampled variance, then it undervalues even more the discretely sampled variance. Our valuations admit enforcement, in some cases, by hedging strategies which perfectly replicate variance swaps by holding log contracts and trading the underlying.

Keywords Variance swap · Lévy process · Time change

Mathematics Subject Classification (2000) 60G51 · 91B28

JEL Classification G13

P. Carr
Courant Institute, NYU, New York, USA
e-mail: pcarr@nyc.rr.com

R. Lee (✉)
Department of Mathematics, University of Chicago, Chicago, USA
e-mail: RL@math.uchicago.edu

L. Wu
Zicklin School of Business, Baruch College, New York, USA
e-mail: liuren.wu@baruch.cuny.edu

Published online: 24 March 2011

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 contracts 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 contracts on stock indices now have bid-offer spreads narrower than those of at-the-money options. The VS has become the standard instrument for taking views on future realized volatility and managing volatility exposure.

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 contracts, the corresponding breakthrough, which we designate as the ND theory (for "Neuberger/Dupire" or "No Discontinuity"), arose in the early 1990s, first in a working paper by Neuberger [15, 16], and then independently in a published article by Dupire [10, 11]. VS contracts began trading sporadically shortly thereafter, and achieved prominence in the late 1990s.

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 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 [1], Dupire [10, 11] first indicated that this forward price can be obtained from co-terminal option prices at all strikes, and Carr and Madan [5] 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 contracts on underlying processes which jump.

The 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 contracts are quoted at twice the estimated forward price of log contracts. Anecdotal evidence and the available historical data confirm that market participants indeed often observe discrepancies between market VS quotes and the ND value.

1.2 Parametric approaches

Notwithstanding the widespread adoption of the nonparametric ND approach, an alternative line of research prices VS contracts 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 et al. [4] find pricing formulas for VS contracts and other volatility derivatives, in terms of the CGMY model's parameters. Under Black–Scholes, Heston, Merton, and Bates dynamics, Broadie and Jain [2] 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 a 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 parametric models allow computation of the (typically small, by [2]) 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 VS pricing.

1.3 Our approach

By introducing jumps in the underlying asset price, we generalize the ND theory of VS pricing. Indeed we value a 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 the special case in which X is Brownian motion. 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, such as the CGMY and Merton models analyzed in some of 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 such as CGMY and Merton. 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 VS pricing. By declining to specify and estimate the dynamics of the clock that generates stochastic variance and jump arrival rates, we decline to price the VS in terms of a full set of estimated parameters. Instead we price the 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. A possible application is to price variance swaps, given vanilla option prices; a modeler who calibrates (imperfectly) a time-changed Lévy process to vanilla option prices can, using our results, discard the calibrated (or *mis*calibrated) time-change parameters, and replace them with the observable price of a log contract. Although the modeler's results will not be robust to erroneous calibration of the Lévy process, they will be robust to erroneous calibration of the time change.

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. A possible application is to calibrate a time-changed Lévy process, given vanilla option prices and given variance swaps; a modeler who observes both log contract prices and variance swap quotes can, by our results, take their ratio to derive an "implied" multiplier and hence an identifying restriction on the parameters of the Lévy process, facilitating the estimation of those parameters. In this setting, a related application is to examine whether time-changed Lévy dynamics prevail in the given market; an implied multiplier that varies significantly across expiries or across time would lead to rejection of all time-changed Lévy dynamics in that market.

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 a Lévy process has a multiplier exceeding 2 if and only if its Lévy measure has negative skewness, in a sense that we define.

Fifth is the capability to enforce our valuations, in some cases, by hedging strategies which perfectly replicate VS payoffs by holding log contracts and trading futures.

Sixth is the extension of nonparametric pricing to discretely sampled VS contracts. In practice, VS contracts specify discrete sampling, but the VS pricing literature mainly addresses continuously sampled variance. An exception is Broadie and Jain [2] which computes discretely sampled VS values, in terms of parameters of four models. We complement this by developing *non*parametric results (for general TCLPs, under an independence condition), including lower bounds on the VS discrete-sampling premium in terms of log contract prices, instead of model parameters. Our lower bounds on this premium are nonnegative. Hence if ND theory undervalues the continuously sampled VS, then in this setting it undervalues even more the discretely sampled VS.

The body 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 gives examples of multiplier formulas; and for some TCLP's calibrated to options data, it computes multipliers which exceed 2 (hence VS prices which exceed ND valuations); and it relates this phenomenon to negative skewness. Section 5 enforces our valuations, in some cases, by hedging strategies which perfectly replicate the VS by holding log contracts and trading futures. Section 6 analyzes the impact of discrete sampling. Section 7 computes the multipliers implied by empirical variance swap data. Section 8 concludes.

2 The multiplier

We work on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_u\}_{u\geq 0}, \mathbb{P})$ satisfying the usual conditions. 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 [19] Propositions 19.2 and 19.5, applied for each $m > 0$ to the restriction of ν to $\{x : 1/m < |x| < m\}$, together with monotone convergence as $m \to \infty$, imply that

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

So we get

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

where the last step is by Sato [19] Corollary 25.8 and the fact that $\int_{|x|<1} x^2 \, d\nu(x) < \infty$. □

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 formula for Q_X 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 equality is by, for instance, Protter [17] Corollary II.6.3. The final formula for Q_X 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*

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

Proof Sato [19] Theorem 25.17 implies that

$$\log \mathbb{E}e^{X_1} = \sigma^2/2 + \int \left(e^x - 1 - x\mathbf{1}_{|x|\le 1}\right)\nu\,(\mathrm{d}x) + \gamma$$

(and that the integral is finite). Sato [19] Example 25.12 implies that

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

(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 function 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 to pay

$$[Y]_T$$

at time 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 [13],

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

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 [14] 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) *Assume* $\mathbb{E}\tau_T < \infty$. *Then*

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

4 Multiplier calculations

In the following examples of returns-driving processes X, we need not 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

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

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 $\mathbb{F} := \{\mathcal{F}_u\}_{u\geq 0}$, an $\mathbb{F}$-Brownian motion W, and a continuous $\mathbb{F}$-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

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

hence

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

is a continuous local martingale. Define the time change $\tau_t := [M]_t = [\log(S_./S_0)]_t$. Let $A_u := \inf\{t : [M]_t \geq u\}$ and $\mathcal{F}_u := \mathcal{G}_{A_u}$. By Dambis [7] and Dubins and Schwarz [9], $W_u := M_{A_u}$ is an $\mathbb{F}$-Brownian motion, and τ is an $\mathbb{F}$-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 and Yor [18] 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) and (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 Time-changed Lévy processes with jumps

In Table 1, we solve for the multipliers of various Lévy processes with jumps. For the background Lévy process, we choose the following examples: a two-jump-size process, the Kou double exponential, the Merton lognormal, the (extended) CGMY, the variance gamma, and the normal inverse Gaussian.

4.3 Impact of skewness

In Table 1, the approximations of Q_X for the two-jump, Kou, Merton, and NIG models exhibit a common theme: increasing *up*-jump sizes (by increasing c_1 or $1/a_1$ or μ or β, respectively) has the leading-order effect of *decreasing* the multiplier, whereas increasing *down*-jump sizes (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 CGMY model, taking larger up-jumps via $(G, M) = (B, b)$, where $B > b$, gives a smaller multiplier than taking larger down-jumps by swapping $(G, M) = (b, B)$. A similar theme emerged in the analysis of the Bates model by Broadie and Jain [2], who found a negative leading-order relationship between the Bates mean jump size parameter and the spread between VS and log contract values.

This asymmetry can be explained as follows. Under any of those dynamics,

$$\begin{aligned}-2\log(F_T/F_0) &= \int_{0+}^{T} \frac{-2}{F_{t-}}\,\mathrm{d}F_t + \frac{1}{2}\int_{0+}^{T} \frac{2}{F_{t-}^2}\,\mathrm{d}[F]_t^c \\ &\quad + \sum_{0<t\le T}\left(-2\Delta\log F_t - \frac{-2}{F_{t-}}\Delta F_t\right) \\ &= \int_{0+}^{T} \frac{-2}{F_{t-}}\,\mathrm{d}F_t + [Y]_T + \sum_{0<t\le T}\left(\frac{2}{F_{t-}}\Delta F_t - 2\Delta Y_t - (\Delta Y_t)^2\right),\end{aligned} \tag{4.1}$$

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

$$[Y]_T + \sum_{0<t\le T}\left(2e^{\Delta Y_t} - 2 - 2\Delta Y_t - (\Delta Y_t)^2\right) \approx [Y]_T + \sum_{0<t\le T}\frac{1}{3}(\Delta Y_t)^3. \tag{4.2}$$

Therefore, in the presence of up-jumps ($\Delta Y_t > 0$), the intuition is that we expect $2\mathbb{E}(-\log F_T/F_0) > \mathbb{E}[Y]_T$, 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 (4.1) and (4.2) resemble closely the jump analysis by Derman et al. [8], 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 (4.2) and Proposition 2.6, we define a relevant notion of skewness.

Table 1 Examples of multipliers for time-changed Lévy processes

Lévy process	Brownian	Lévy density	Multiplier (exact)	Multiplier (approximate)
Brownian	σ^2	0	2	2
Two jump sizes[a]	σ^2	$\lambda_1\delta_{c_1}+\lambda_2\delta_{c_2}$	$\frac{\sigma^2+\lambda_1 c_1^2+\lambda_2 c_2^2}{\sigma^2/2+\lambda_1(e^{c_1}-1-c_1)+\lambda_2(e^{c_2}-1-c_2)}$	$2-\frac{2\lambda_1}{3\sigma^2}c_1^3+\frac{2\lambda_2}{3\sigma^2}\lvert c_2\rvert^3$
Kou double exp[b]	σ^2	$\lambda_1 a_1 e^{-a_1\lvert x\rvert}\mathbf{1}_{x>0}+\lambda_2 a_2 e^{-a_2\lvert x\rvert}\mathbf{1}_{x<0}$	$\frac{\sigma^2+2\lambda_1/a_1^2+2\lambda_2/a_2^2}{\sigma^2/2+\lambda_1/(a_1-1)-\lambda_2/(a_2+1)-\lambda_1/a_1+\lambda_2/a_2}$	$2-\frac{4\lambda_1/\sigma^2}{a_1^3}+\frac{4\lambda_2/\sigma^2}{a_2^3}$
Merton[c]	σ^2	$\frac{\lambda}{\eta\sqrt{2\pi}}\exp(\frac{-(x-\mu)^2}{2\eta^2})$	$\frac{\sigma^2+\lambda\eta^2+\lambda\mu^2}{\sigma^2/2+\lambda(e^{\mu+\eta^2/2}-1-\mu)}$	$2-\frac{2\lambda}{\sigma^2}\eta^2\mu-\frac{2\lambda}{3\sigma^2}\mu^3$
General CGMY[d]	0	$\frac{C_n}{\lvert x\rvert^{1+Y_n}}e^{-G\lvert x\rvert}\mathbf{1}_{x<0}+\frac{C_p}{\lvert x\rvert^{1+Y_p}}e^{-M\lvert x\rvert}\mathbf{1}_{x>0}$	Q	$\tilde{Q}$
CGMY[e]	0	$\frac{C}{\lvert x\rvert^{1+Y}}(e^{-G\lvert x\rvert}\mathbf{1}_{x<0}+e^{-M\lvert x\rvert}\mathbf{1}_{x>0})$	$\frac{Y(1-Y)(G^{Y-2}+M^{Y-2})}{G^Y-(G+1)^Y+YG^{Y-1}+M^Y-(M-1)^Y-YM^{Y-1}}$	$\frac{2G^{Y-2}+2M^{Y-2}}{G^{Y-2}(1-\frac{2-Y}{3G}+\cdots)+M^{Y-2}(1+\frac{2-Y}{3M}+\cdots)}$
VG[f]	0	$\frac{C}{\lvert x\rvert}(e^{-G\lvert x\rvert}\mathbf{1}_{x<0}+e^{-M\lvert x\rvert}\mathbf{1}_{x>0})$	$\frac{1/G^2+1/M^2}{1/G-\log(1+1/G)-1/M-\log(1-1/M)}$	$\frac{2G^2+2M^2}{M^2(1-\frac{2}{3G}+\cdots)+G^2(1+\frac{2}{3M}+\cdots)}$
NIG[g]	0	$\frac{\delta\alpha}{\pi}\frac{\exp(\beta x)K_1(\alpha\lvert x\rvert)}{\lvert x\rvert}$	$\frac{\alpha^2/(\alpha^2-\beta^2)}{\alpha^2-\beta^2-\beta-\sqrt{(\alpha^2-\beta^2)(\alpha^2-(\beta+1)^2)}}$	$2-\frac{4\beta+1}{2\alpha^2}$

[a]Let δ_c denote 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 $\lvert c_2\rvert$. The multiplier approximation is by a third-order Taylor expansion in (c_1, c_2) about $(0,0)$

[b]Let $a_1 \geq 1$ and $a_2>0$; hence up-jumps and down-jumps have mean absolute size $1/a_1$ and $1/a_2$, respectively. The multiplier approximation is by a third-order Taylor expansion in $(1/a_1, 1/a_2)$ about $(0,0)$, provided $\sigma \neq 0$. 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 case, the multiplier becomes exactly $2-2/a_1+2/a_2$

[c]The multiplier approximation is by a third-order Taylor expansion in (μ, η) about $(0,0)$, provided that $\sigma \neq 0$

[d]Here $Q := \frac{C_n\Gamma(2-Y_n)G^{Y_n-2}+C_p\Gamma(2-Y_p)M^{Y_p-2}}{C_n\Gamma(-Y_n)[(G+1)^{Y_n}-G^{Y_n}-Y_nG^{Y_n-1}]+C_p\Gamma(-Y_p)[(M-1)^{Y_p}-M^{Y_p}+Y_pM^{Y_p-1}]}$ and $\tilde{Q} := \frac{2C_n\Gamma(2-Y_n)G^{Y_n-2}+2C_p\Gamma(2-Y_p)M^{Y_p-2}}{C_n\Gamma(2-Y_n)G^{Y_n-2}(1-\frac{2-Y_n}{3G}+\cdots)+C_p\Gamma(2-Y_p)M^{Y_p-2}(1+\frac{2-Y_p}{3M}+\cdots)}$, where $C_p, C_n, G>0$, and $M>1$, and $Y_p, Y_n<2$, with $\{Y_p, Y_n\}\cap\{0,1\}=\emptyset$. The approximation $\tilde{Q}$ is by expansion, in $1/G$ and $1/M$, of the denominator of Q

[e]Let $C, G>0$, and $M>1$, and $Y<2$

[f]Let $C, G>0$ and $M>1$

[g]Let $\delta>0$, $\alpha>0$, and $-\alpha<\beta<\alpha-1$. Let K_1 denote the modified Bessel function of the second kind and order 1. 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

Definition 4.2 (Exponential skewness) For a Lévy measure ν such that

$$\int_{|x|>1} e^x \nu(\mathrm{d}x) < \infty \quad \text{and} \quad \int_{|x|>1} x^2 \nu(\mathrm{d}x) < \infty,$$

define the *exponential skewness* of ν by

$$6 \int \left(e^x - 1 - x - x^2/2\right) \nu(\mathrm{d}x). \tag{4.3}$$

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 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 (4.3), exponential skewness is negative if and only if

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

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

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

4.4 Multipliers of empirically calibrated processes

Carr et al. [3] calibrate various time-changed Lévy processes to data. In Table 2, 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 et al. 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.

5 Perfect hedging

In some cases, our valuation results are enforceable (assuming frictionless markets) by perfect hedging strategies which hold log contracts statically and trade futures dynamically.

Table 2 Carr et al. [3] calibration, using 4 cross-sections of S&P500 options data in 2000

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

5.1 One jump size

Consider the case that X has one possible jump size $c \neq 0$, and zero Brownian part. In other words, X is c times a simple Poisson process with drift. By Proposition 2.6, the multiple

$$Q_X = \frac{c^2}{e^c - 1 - c}$$

of a log contract prices the variance swap. Proposition 5.1 shows, moreover, that this valuation is enforceable by the following hedging strategy: Hold Q_X log contracts statically, together with $e^{R_t - R_T} Q_X / F_{t-}$ futures dynamically (and storing the resulting profits/losses in $\int_0^t Q_X/F_{u-}\, \mathrm{d}F_u$ bonds) for each $t \in (0, T)$, producing a final portfolio value equal to the variance payoff $[Y]_T$.

Proposition 5.1 *Let X have zero Brownian part and Lévy measure $\nu = \delta_c$, where $c \in \mathbb{R} \setminus \{0\}$. Then*

$$Q_X \log(F_0/F_T) + \int_0^T \frac{Q_X}{F_{t-}}\, \mathrm{d}F_t = [Y]_T. \tag{5.1}$$

Proof We are given

$$\bar{X}_u = mu + cN_u,$$

where N is a Poisson process and m is a constant that we need not specify. Re-indexing by calendar time, we have

$$Y_t = m\tau_t + c\tilde{N}_t, \tag{5.2}$$

where $\tilde{N}_t := N_{\tau_t}$. By Itô's rule, the futures price $F_t = F_0 \exp(Y_t)$ satisfies

$$\mathrm{d}F_t = mF_{t-}\, \mathrm{d}\tau_t + \big(e^c - 1\big) F_{t-}\, \mathrm{d}\tilde{N}_t. \tag{5.3}$$

Combining (5.2) and (5.3),

$$-Q_X \,\mathrm{d}\log F_t + \frac{Q_X}{F_{t-}}\,\mathrm{d}F_t = Q_X\big(e^c - c - 1\big)\,\mathrm{d}\tilde{N}_t = c^2\,\mathrm{d}\tilde{N}_t,$$

which implies (5.1). □

5.2 Two jump sizes, piecewise constant paths

Consider the case that X has two possible jump sizes $c_1 > 0$ and $c_2 < 0$, zero Brownian part, and piecewise constant paths. In other words, X has Lévy measure

$$p_1\delta_{c_1} + p_2\delta_{c_2}, \tag{5.4}$$

where

$$p_1 := \frac{1-e^{c_2}}{e^{c_1}-e^{c_2}}, \qquad p_2 := \frac{e^{c_1}-1}{e^{c_1}-e^{c_2}}, \tag{5.5}$$

producing paths (of X, Y, and F) which are constant between jumps. This means that X is a linear combination of two simple Poisson processes with drift.

By Proposition 2.6, the multiple

$$Q_X = \frac{p_1c_1^2 + p_2c_2^2}{p_1(e^{c_1}-1-c_1) + p_2(e^{c_2}-1-c_2)} = \frac{c_1^2(e^{c_2}-1) - c_2^2(e^{c_1}-1)}{c_2(e^{c_1}-1) - c_1(e^{c_2}-1)}$$

of a log contract prices the variance swap. Proposition 5.2 shows, moreover, that this valuation is enforceable by the following hedging strategy: Hold Q_X log contracts statically, together with $e^{R_t-R_T}q_X/F_{t-}$ futures dynamically (and storing the resulting profits/losses in $\int_0^t q_X/F_{u-}\,\mathrm{d}F_u$ bonds) for each $t \in (0,T)$, where

$$q_X := \frac{c_1c_2(c_1-c_2)}{c_2(e^{c_1}-1) - c_1(e^{c_2}-1)},$$

producing a final portfolio value equal to the variance payoff $[Y]_T$.

Proposition 5.2 *Let X have zero Brownian part and Lévy measure* (5.4), (5.5). *Then*

$$Q_X\log(F_0/F_T) + \int_0^T \frac{q_X}{F_{t-}}\,\mathrm{d}F_t = [Y]_T. \tag{5.6}$$

Proof We have

$$\bar{X}_u = c_1N_u^1 + c_2N_u^2,$$

where N^1 and N^2 are independent Poisson processes. Re-indexing by calendar time, we have

$$Y_t = c_1\tilde{N}_t^1 + c_2\tilde{N}_t^2, \tag{5.7}$$

where $\tilde{N}_t^j := N_{\tau_t}^j$ for $j = 1, 2$. By Itô's rule, the futures' price $F_t = F_0 \exp(Y_t)$ satisfies

$$\mathrm{d}F_t = \big(e^{c_1} - 1\big) F_{t-}\, \mathrm{d}\tilde{N}_t^1 + \big(e^{c_2} - 1\big) F_{t-}\, \mathrm{d}\tilde{N}_t^2. \tag{5.8}$$

Combining (5.7) and (5.8),

$$-Q_X \,\mathrm{d}\log F_t + \frac{q_X}{F_{t-}}\,\mathrm{d}F_t = \sum_{j=1,2} \big(q_X e^{c_j} - q_X - Q_X c_j\big)\, \mathrm{d}\tilde{N}_t^j = c_1^2\, \mathrm{d}\tilde{N}_t^1 + c_2^2\, \mathrm{d}\tilde{N}_t^2,$$

which implies (5.6). □

6 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 (unannualized) payoffs of (the floating leg of) a *discretely sampled variance swap* on futures F and on spot F^* to be, respectively,

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

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

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

Still working under Sect. 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 6.1 (Discrete variance swap on spot underlying: Valuation) *Assume that* $\mathbb{E}\tau_T < \infty$ *and that* τ *and* X *are independent. Then*

$$\mathbb{E}V_T^* = \mathbb{E}[Y^*]_T + \sum_{n=0}^{N-1} (\mathbb{E}\Delta_n Y^*)^2 + \sum_{n=0}^{N-1} \mathrm{Var}\big(\mathbb{E}(\Delta_n Y^* | \tau)\big). \tag{6.1}$$

The last term has the explicit form

$$\sum_{n=0}^{N-1} \operatorname{Var}\big(\mathbb{E}(\Delta_n Y^*|\tau)\big) = (\mathbb{E}X_1)^2\, \mathbb{E}\sum_{n=0}^{N-1}(\Delta_n \tau)^2 - \sum_{n=0}^{N-1}(\mathbb{E}\Delta_n Y)^2. \tag{6.2}$$

Proof By the finiteness of $\mathbb{E}\tau_T$ and the Lévy, strong Markov and martingale properties of

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

Wald's first equation implies that for $t \in [0, T]$,

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

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

$$\mathbb{E}\big(\Delta[Y^*]\big) = \mathbb{E}\big(\Delta[Y]\big) = \mathbb{E}\big(\Delta[M]\big) = \mathbb{E}(\Delta M)^2 = \mathbb{E}\big(\Delta Y - (\Delta\tau)\mathbb{E}X_1\big)^2, \tag{6.3}$$

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 [17] Corollary II.6.3 on p. 73. By the independence condition,

$$\mathbb{E}(\Delta Y|\tau) = (\Delta\tau)\mathbb{E}X_1, \tag{6.4}$$

which implies that (6.3) becomes

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

by the nonrandomness of $Y - Y^*$. Summing from $n = 0$ to $N - 1$ proves (6.1). By (6.4), $\operatorname{Var}(\mathbb{E}(\Delta Y|\tau)) = \mathbb{E}(\Delta\tau\mathbb{E}X_1)^2 - (\mathbb{E}\Delta Y)^2$, which proves (6.2). □

Corollary 6.2 (Discrete variance swap on spot underlying: Lower bound) *Under the assumptions of Proposition* 6.1, *we have*

$$\mathbb{E}V_T^* \geq \mathbb{E}[Y^*]_T + \sum_{n=0}^{N-1}(\mathbb{E}Y^*_{t_{n+1}} - \mathbb{E}Y^*_{t_n})^2. \tag{6.5}$$

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 (6.1) and $\operatorname{Var}(\mathbb{E}(\Delta Y^*|\tau)) \geq 0$, with equality if τ is nonrandom. □

Proposition 6.3 (Discrete variance swap on futures) *By deleting all instances of stars* (*) *in their statements, Proposition* 6.1 *and Corollary* 6.2 *apply to discretely sampled variance swaps on futures.*

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

Under the Black–Scholes and Merton jump-diffusion models, Broadie and Jain [2] found that discrete sampling does theoretically increase variance swap values. We regard those models as instances of exponential Lévy processes under a nonrandom clock; hence (6.5) holds with equality, thereby expressing the discrete-sampling premium in terms of log contract prices.

More generally, Corollary 6.2 implies that for general exponential Lévy processes time-changed by independent stochastic clocks, the discrete-sampling premium $\mathbb{E}V_T^* - \mathbb{E}[Y^*]_T$ is still nonnegative and bounded below in terms of log contract prices. So if the commonly quoted multiple 2 undervalues the continuously sampled variance swap (as suggested by the data in Sect. 4.4), then in this setting, the multiple 2 furthermore undervalues the discretely sampled variance swap.

7 Multiplier estimates from S&P variance swap data

Whereas Sect. 4.4 estimated multipliers from empirically calibrated parameters of the Lévy measure, this section estimates multipliers from empirical observations of variance swap quotes and log contract valuations, by taking the ratio of the former and the latter. As suggested by a referee, nonconstancy of this ratio in empirical data would be evidence against the family of time-changed Lévy processes for modeling the dynamics underlying that data set.

From a major broker-dealer, we obtained daily closing quotes on variance swaps on the S&P500 index with fixed times to expiry of 2, 3, 6, 12, and 24 months. To avoid the effect of weekday patterns on the dynamics estimation, we sampled the data weekly on every Wednesday. When Wednesday was a holiday, we used the most recent observation before the holiday. The data contain 682 weekly observations for each series, from January 10, 1996 to January 28, 2009.

To construct log contract valuations at the corresponding dates and times to expiry, we use option valuation data, expressed in terms of implied volatility. Specifically, we retrieve the daily closing bid and ask quotes on the S&P500 index options from OptionMetrics, and we apply the following filters. First, we retain only options quotes with strictly positive bids and with times to expiry of no fewer than seven calendar days. Second, we retain only the expiries for which we can obtain valid implied volatilities for at least 10 strikes at that expiry, where our implied volatility calculation uses option values obtained from the mid-market price at each available strike, discount factors obtained from the Eurodollar LIBOR and swap rates (via Bloomberg), and the underlying forward price obtained as the option-implied forward index level which best reconciles call and put prices via put-call parity.

The above filtering yields 169,669 implied volatilities over the 682 days of our sample period. Each date contains three to ten expiries, with an average of about eight.

The expiries range from the minimum requirement of seven days up to about three years. Within each date and expiry, the number of strikes ranges from the minimum requirement of 10 up to 150, with an average of 31 strikes per expiry.

To value the log contract for each date t and expiry T, we integrate squared implied volatility IV^2, regarded as a function of log-strike $k = \log(K/F_t)$ across all $k \in \mathbb{R}$, with weightings given by the normal distribution function Φ, evaluated at the standardized log-strike $m(k)$, i.e.,

$$\mathbb{E}_t \log(F_t/F_T) = \frac{T-t}{2} \int_{-\infty}^{\infty} IV^2(k)\, \mathrm{d}\Phi\big(m(k)\big),$$
$$m(k) := \frac{k}{IV(k)\sqrt{T-t}} + \frac{IV(k)\sqrt{T-t}}{2}. \tag{7.1}$$

See Gatheral [12, Chap. 11, (11.5)] for a derivation of this representation of the log contract valuation.

To evaluate (7.1) numerically, we interpolate and extrapolate from the available strike data, producing IV estimates at all log-strikes k on a fine and extensive grid. Specifically, we use linear interpolation and flat extrapolation, as used also in [6]. The interpolation generates IV^2 linearly between all pairs of adjacent available strikes. The flat extrapolation assigns IV^2 at all strikes below the lowest (respectively, above the highest) available strike to be equal to the IV^2 at the lowest (respectively, highest) available strike. Results from three alternative schemes, including a sloping extrapolation, are qualitatively similar and are available upon request.

After computing the log contract valuation at each option expiry T, we perform linear (in T) interpolation on the expected-variance term structure $T \mapsto \mathbb{E}_t \log(F_t/F_T)$, to obtain the log contract valuation at times to expiry of 2, 3, 6, 12, and 24 months, corresponding to the variance swap data.

We calculate the ratio of the variance swap quote to the log contract value. By Proposition 3.1, this "implied multiplier" should be constant across times and expiries if the underlying dynamics belong to the family of all time-changed Lévy processes. In Table 3 and Fig. 1, however, we observe variations of the ratio, which was mostly greater than 2.0 in the early part of the sample, but turned smaller than 2.0 in the more recent part of the sample. These variations suggest that time-changed Lévy processes of the form (3.1) may not fully describe the dynamics of S&P500 returns.

8 Conclusion

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 widely quoted volatility indicators such as the VIX, VXN, and VSTOXX. 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,

Table 3 Ratio of variance swap quote to log contract value for S&P500

Subsample		Time to expiry (months)				
		2	3	6	12	24
Full sample	Mean	2.01	1.99	2.02	2.07	2.10
$N = 682$	St dev	0.09	0.09	0.09	0.11	0.13
1997–1999	Mean	2.06	2.07	2.10	2.17	2.20
$N = 156$	St dev	0.09	0.08	0.08	0.10	0.09
2000–2002	Mean	2.05	2.01	2.06	2.15	2.19
$N = 157$	St dev	0.06	0.05	0.05	0.06	0.08
2003–2005	Mean	1.97	1.93	1.93	1.98	2.01
$N = 156$	St dev	0.06	0.05	0.05	0.06	0.08
2006–2008	Mean	1.96	1.95	1.95	1.97	1.96
$N = 157$	St dev	0.06	0.06	0.04	0.03	0.05

The full sample runs from 2006 January 10 to 2009 January 28. Subsampling of three-year periods shows ratios greater than 2.0 in the early part, but smaller than 2.0 in the later part

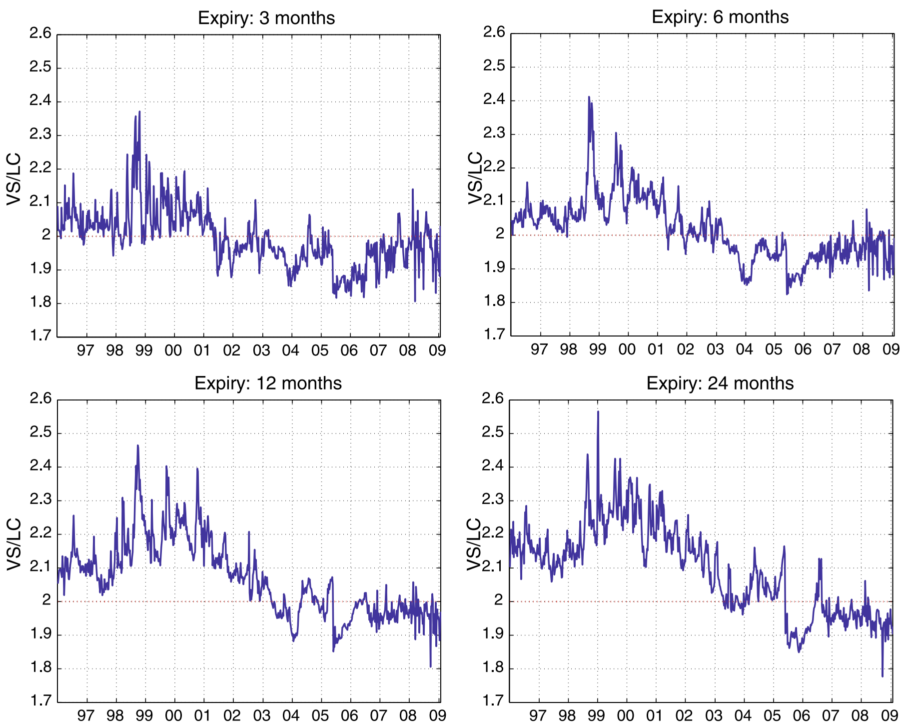

Fig. 1 Ratio of variance swap quote to log contract value for S&P500. The *solid line* in each panel plots the time series of VS/LC, the ratio of the variance swap quote to the log contract value, for the S&P500. In the absence of jump risk, this ratio would be 2.0

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 exponential 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.

We show that in some cases of one or two possible jump sizes, our valuations admit enforcement by hedging strategies which perfectly replicate the variance swap payoff by holding log contracts statically and trading the underlying dynamically.

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

Finally we compute the ratio of variance swap quotes to log contract valuations in S&P500 data, and we observe variations in this ratio, across time and expiry.

Acknowledgements We thank Tom Bielecki and two anonymous referees for helpful comments.

References

1. Breeden, D., Litzenberger, R.: Prices of state contingent claims implicit in options prices. J. Bus. **51**, 621–651 (1978)
2. Broadie, M., Jain, A.: The effect of jumps and discrete sampling on volatility and variance swaps. Int. J. Theor. Appl. Finance **11**, 761–797 (2008)
3. Carr, P., Geman, H., Madan, D., Yor, M.: Stochastic volatility for Lévy processes. Math. Finance **13**, 345–382 (2003)
4. Carr, P., Geman, H., Madan, D., Yor, M.: Pricing options on realized variance. Finance Stoch. **9**, 453–475 (2005)
5. Carr, P., Madan, D.: Towards a theory of volatility trading. In: Jarrow, R. (ed.) Volatility, pp. 417–427. Risk Publications, London (1998)
6. Carr, P., Wu, L.: Variance risk premia. Rev. Financ. Stud. **22**, 1311–1341 (2009)
7. Dambis, K.E.: On the decomposition of continuous submartingales. Theory Probab. Appl. **10**, 401–410 (1965)
8. Derman, E., Demeterfi, K., Kamal, M., Zou, J.: A guide to volatility and variance swaps. J. Deriv. **6**, 9–32 (1999)

9. Dubins, L.E., Schwarz, G.: On continuous martingales. Proc. Natl. Acad. Sci. USA **53**, 913–916 (1965)
10. Dupire, B.: Arbitrage Pricing with Stochastic Volatility. Société Générale, Paris (1992)
11. Dupire, B.: Model art. Risk **6**(9), 118–124 (1993)
12. Gatheral, J.: The Volatility Surface: A Practitioner's Guide. Wiley, New York (2006)
13. Hall, W.J.: On Wald's equations in continuous time. J. Appl. Probab. **7**, 59–68 (1970)
14. Jacod, J.: Calcul Stochastique et Problèmes de Martingales. Lecture Notes in Mathematics, vol. 714. Springer, Berlin (1979)
15. Neuberger, A.: Volatility trading. London Business School, Working Paper (1990)
16. Neuberger, A.: The log contract. J. Portf. Manag. **20**, 74–80 (1994)
17. Protter, P.: Stochastic Integration and Differential Equations, 2nd edn. Springer, Berlin (2004)
18. Revuz, D., Yor, M.: Continuous Martingales and Brownian Motion, 3rd edn. Springer, Berlin (1999)
19. Sato, K.-I.: Lévy Processes and Infinitely Divisible Distributions. Cambridge University Press, Cambridge (1999)