Variance Swap: Replication and Hedging

Guanlin Li

Contents

1 Realized Variance

1.1 Setup

Let F_t denote the futures price. Under the futures measure and with zero carry for the futures itself,

$$dF_t = \sigma_t F_t \, dW_t. \tag{1}$$

Equivalently,

$$\frac{dF_t}{F_t} = \sigma_t dW_t. \tag{2}$$

The continuously sampled realized variance over $[0, T]$ is

$$RV_{0,T} := \int_0^T \sigma_t^2 \, dt. \tag{3}$$

1.2 Realized Variance Replication: Ito's lemma

Consider the log contract

$$g(F_t) = -2 \log \left(\frac{F_t}{F_0} \right). \tag{4}$$

Its first and second derivatives are

$$g'(F) = -\frac{2}{F}, \qquad g''(F) = \frac{2}{F^2}. \tag{5}$$

Applying Ito's lemma,

$$dg(F_t) = g'(F_t) dF_t + \frac{1}{2} g''(F_t)(dF_t)^2. \tag{6}$$

Since $(dF_t)^2 = \sigma_t^2 F_t^2 dt$, we obtain

$$dg(F_t) = -\frac{2}{F_t} dF_t + \frac{1}{2} \cdot \frac{2}{F_t^2} \cdot \sigma_t^2 F_t^2 dt \tag{7}$$

$$= -\frac{2}{F_t} dF_t + \sigma_t^2 dt. \tag{8}$$

Rearranging,

$$\sigma_t^2 dt = dg(F_t) + \frac{2}{F_t} dF_t. \tag{9}$$

Integrating from 0 to T,

$$\int_0^T \sigma_t^2 dt = -2 \log\left(\frac{F_T}{F_0}\right) + \int_0^T \frac{2}{F_t} dF_t. \tag{10}$$

Hence the realized variance payoff admits the replication

$$RV_{0,T} = \underbrace{-2 \log\left(\frac{F_T}{F_0}\right)}_{\text{static log contract}} + \underbrace{\int_0^T \frac{2}{F_t} dF_t}_{\text{dynamic futures position}}. \tag{11}$$

The second term means simultaneously holding $2/F_t$ of futures at any time t. If F_t goes up, the position needs to be rebalanced by selling (as $1/F_t$ decreases), vice versa. The rebalancing trades, buy low and sell high, implies long convexity.

1.3 Variance Scaling and Time Convention

To avoid ambiguity in variance scaling, we measure time in days (say $dt = 1$). Under this convention, σ_t represents the daily volatility (not annualized), its square (variance) is defined as $\sigma_t^2 = \mathbb{E}[(\log(F_t/F_{t-1}))^2]$.

1.4 First and Second Derivatives of the Log Contract

For

$$g(F) = -2 \log\left(\frac{F}{F_0}\right), \tag{12}$$

the delta and gamma are

$$\Delta_g(F) = g'(F) = -\frac{2}{F}, \qquad \Gamma_g(F) = g''(F) = \frac{2}{F^2}. \tag{13}$$

Therefore the cash gamma is

$$F^2 \Gamma_g(F) = F^2 \cdot \frac{2}{F^2} = 2, \tag{14}$$

which is constant. The log contract has *constant cash gamma.* In particular, if the log contract is delta hedged continuously, then the local second-order (convexity) P&L contribution is

$$\text{Gamma P\&L} = \frac{1}{2} \Gamma_g(F_t)(dF_t)^2 = \frac{1}{2} \cdot \frac{2}{F_t^2} (dF_t)^2 = \left(\frac{dF_t}{F_t}\right)^2, \tag{15}$$

which is exactly the realized variance contribution.

2 Static Replication of Log Contract

2.1 Portfolio of Options

Carr–Madan static replication gives, for a twice differentiable payoff $h(F_T)$ expanded around F_0,

$$h(F_T) = h(F_0) + h'(F_0)(F_T - F_0) \tag{16}$$

$$+ \int_0^{F_0} h''(K)(K - F_T)^+ dK + \int_{F_0}^{\infty} h''(K)(F_T - K)^+ dK. \tag{17}$$

Setting

$$h(F_T) = -2\log\left(\frac{F_T}{F_0}\right), \tag{18}$$

we have

$$h(F_0) = 0, \qquad h'(F_0) = -\frac{2}{F_0}, \qquad h''(K) = \frac{2}{K^2}. \tag{19}$$

Hence, $g(F_T)$ is replicated by follows,

$$-2\log\left(\frac{F_T}{F_0}\right) = -\frac{2}{F_0}(F_T - F_0) + \int_0^{F_0} \frac{2}{K^2}(K - F_T)^+ dK + \int_{F_0}^{\infty} \frac{2}{K^2}(F_T - K)^+ dK. \tag{20}$$

So the static log contract can be replaced by:

- one linear futures leg with quantity $-2/F_0$,
- a strip of OTM puts for $K < F_0$ weighted by $2/K^2$,
- a strip of OTM calls for $K > F_0$ weighted by $2/K^2$.

Substituting this into the previous replication gives

$$RV_{0,T} = \int_0^T \frac{2}{F_t} dF_t - \frac{2}{F_0}(F_T - F_0) + \int_0^{F_0} \frac{2}{K^2}(K - F_T)^+ dK + \int_{F_0}^{\infty} \frac{2}{K^2}(F_T - K)^+ dK. \tag{21}$$

$$= \underbrace{\int_0^T \left(\frac{2}{F_t} - \frac{2}{F_0}\right) dF_t}_{\text{holding } 1/F_t - 1/F_0 \text{ of futures at time } t} + \underbrace{\int_0^{F_0} \frac{2}{K^2}(K - F_T)^+ dK + \int_{F_0}^{\infty} \frac{2}{K^2}(F_T - K)^+ dK.}_{\text{holding a portfolio of calls and puts, weighted by } 1/K^2} \tag{22}$$

Note that in a Black-Scholes framework, the first part, *i.e.,* simultaneously holding $1/F_t - 1/F_0$ of futures at time t, continuously rebalanced futures holding encapsulates the delta-hedging of the options [AEG06]. Therefore, the forward cost of this portfolio (taking expectation) is simply:

$$K^2_{\text{VAR}} = 2e^{rT}\left[\int_0^{F_0} \frac{P_0(K)}{K^2} dK + \int_{F_0}^{\infty} \frac{C_0(K)}{K^2} dK.\right] \tag{23}$$

2.2 Option Strip Delta

The static option strip is the pure convexity part of the log-contract decomposition. Define

$$\text{Strip}(F_T) := \int_0^{F_0} \frac{2}{K^2}(K - F_T)^+ dK + \int_{F_0}^{\infty} \frac{2}{K^2}(F_T - K)^+ dK. \tag{24}$$

From the Carr–Madan decomposition,

$$-2\log\left(\frac{F_T}{F_0}\right)=-\frac{2}{F_0}(F_T-F_0)+\mathrm{Strip}(F_T). \tag{25}$$

Note that the above relationship should remain valid up to expiration T (arbitrage-free assumption). Differentiating both sides with respect to F_T gives

$$-\frac{2}{F_T}=-\frac{2}{F_0}+\Delta_{\mathrm{strip}}(F_T), \tag{26}$$

so the strip delta is

$$\Delta_{\mathrm{strip}}(F_T)=-\frac{2}{F_T}+\frac{2}{F_0}. \tag{27}$$

Evaluating at the expansion or initial point $F_T=F_0$,

$$\Delta_{\mathrm{strip}}(F_0)=-\frac{2}{F_0}+\frac{2}{F_0}=0. \tag{28}$$

Why the Option Strip Delta Equals Zero at F_0 ? This result has a clear interpretation. The first-order exposure of the log contract at the anchor point F_0 is carried entirely by the linear term $-\frac{2}{F_0}(F_T-F_0)$. The option strip only carries the higher-order curvature needed to bend the payoff away from that tangent line. Therefore, at the anchor point F_0, the strip has zero delta by construction. The strip does not carry any initial directional exposure at the forward anchor, but it does carry the convexity required for constant cash gamma.

2.3 Variance Swap Valuation and Daily Rebalancing

The value of a variance swap, at time t is given by (written at $t=0$, expiry at $t=T$, where $0\le t\le T$):

$$P_t=M\times\left[RV_{0,T}^{\mathrm{Expected}}(t)-K_0^2\right], \tag{29}$$

where M is some notional or scaling constant (omitted), K_0 is the initial fixed variance swap strike and $RV_{0,T}^{\mathrm{Expected}}(t)$ is the expected realized variance up to time t. This expected variance is then calculated in the same way as the mark-to-market (mtm) P&L as a combination accrued realized variance to date and future implied variance [AEG06].

Note that intra-day there is a term representing the square of the move which will act to give the variance swap delta on an intra-day basis, assume $t'\le t\le t'+1$,

$$RV_{0,T}^{\mathrm{Expected}}(t)=\underbrace{RV_{0,t'}(\mathcal{F}_t)}_{\text{realized variance(observed)}}\underbrace{+(t-t')\left[\log\left(\frac{F_t}{F_{t'}}\right)^2\right]}_{\text{intraday variance}}+\underbrace{K_{t,T}^2}_{\text{implied variance}}, \tag{30}$$

where t' is the close time of last day, $RV_{0,t'}$ is the realised variance accrued to last day, $K_{t,T}^2$ is the variance strike (centered at F_t) evaluated at day time t expiring at T. $F_{t'}$ is the value of underlying at the close on last day, and F_t is the value of the underlying at the valuation time on day time t. At the end of day t, the intra-day variance will be absorbed to realized variance component. The Eq. 30 differs from the decomposition presented in Sec. 4.7(*Variance swap Greeks*) of [AEG06], this discrepancy is solely due to our scaling convention, as detailed in Sec. 1.3.

The variance swap can be seen to take on delta only intra-day, we calculate $\partial RV_{0,T}^{\mathrm{Expected}}(t)/\partial F_t$, assume t is the close time of day t, so $t-t'=1$, then,

$$\frac{\partial RV_{0,T}^{\mathrm{Expected}}(t)}{\partial F_t}=\frac{2}{F_t}\log\left(\frac{F_t}{F_{t-1}}\right)\approx\frac{2}{F_t}\left(\frac{F_t}{F_{t-1}}-1\right). \tag{31}$$

Recall Eq. 21, the first term represents the daily holding position. We now compute the rebalancing delta, defined as the difference between the position at day t and day $t-1$, which exactly offsets the intra-day delta exposure discussed above:

$$\left(\frac{2}{F_t} - \frac{2}{F_0}\right) - \left(\frac{2}{F_{t-1}} - \frac{2}{F_0}\right) = -\frac{2}{F_t}\left(\frac{F_t}{F_{t-1}} - 1\right). \tag{32}$$

This represents the replication of the log contract which will have to be done at the end of the day to capture that day's realised variance. Note that, the implied variance term in Eq. 30, may depend on the underlying (linked by volatility skew, e.g., negative skew implies that underlying rallies then the implied variance goes down), then the variance swap acquires other sources of delta, in addition to the intra-day delta discussed above. Similarly, vega exposure should be understood as sensitivity to the shape of the volatility surface (*ATM, skew, kurtosis*, etc.), not just a single-strike point [AEG06].

'Counter-intuitive' nature of the hedge position. A somewhat counter-intuitive feature of variance swap replication is that the *target* hedge (holding) position is completely independent of the volatility level (see Eq. 21). At first glance, one might therefore conclude that the rebalancing effort is invariant to market conditions? In practice, however, both the *volume* of daily trades and the resulting P&L are highly sensitive to volatility. To see this, note that the daily adjustment in the hedge position is

$$d\left(\frac{2}{F_t}\right) = \frac{2}{F_t}\frac{dF_t}{F_t}, \tag{33}$$

which scales directly with the magnitude of the price move dF_t/F_t. Hence, in a high-volatility environment, where dF_t/F_t is larger, the trader must buy or sell substantially more shares each day. The dynamic rebalancing process automatically extracts a P&L that is proportional to the realised variance making the strategy's profitability entirely driven by the realised volatility.

2.4 Hedging Spot Variance Swaps with Futures: The Spot-Basis Correlation

A market-maker who is long the index (spot) variance swap (*i.e.,* client is a variance swap seller, *Sell Vol*) can offset the risk by shorting the replicating option strip and delta-hedging. They will therefore be short gamma in the options market. This short gamma exposure is the hedge against the client's short-volatility position; **while the overall swap-plus-replication portfolio is designed to be gamma-neutral**, the market-maker deliberately maintains a short gamma exposure in the options leg of the replication to offset the long variance (swap) exposure inherited from the swap deal. A delta-hedger who is short options will act to buy underlying as it rallies, and selling as it sells-off. Therefore, this market maker will receive the floating realised variance from the client and pay the fixed strike K_{var}.

Now, suppose the market maker executes the delta-hedge using futures contracts instead of spot instruments, driven by liquidity constraints and regulatory policies specific to the Chinese market. Let $B_t = F_t - S_t$ denote the basis, where F_t is the futures price and S_t is the spot price. In the presence of a systematic positive correlation between the spot and the basis, where a sell-off deepens the backwardation (i.e., B_t goes down when S_t drops) and a rally strengthens the contango (i.e., B_t goes up when S_t rises). Consequently, Variance(F_t) > Variance(S_t), where $F_t = S_t + B_t$.This positive spot-basis correlation inflates the effective cost of delta-hedging relative to the variance payoff expected from the client (which is settled using close-to-close index spot variance). As a result, the market maker incurs larger hedging losses than the realised variance they receive, effectively paying out more variance through the hedge than they collect from the variance swap seller.

3 Numerical Illustration

Figure 1 shows the numerical simulation and decomposition generated in the accompanying Python code. The plots compare:

- realized variance against the exact log-contract replication, see Eq. 10
- the option-strip based replication and its hedge P&L, see Eq. 21
- the theoretical hedge by holding $(2/F_t - 2/F_0)$ futures against the numerically computed option strip delta (negative signed), see Eq. 27

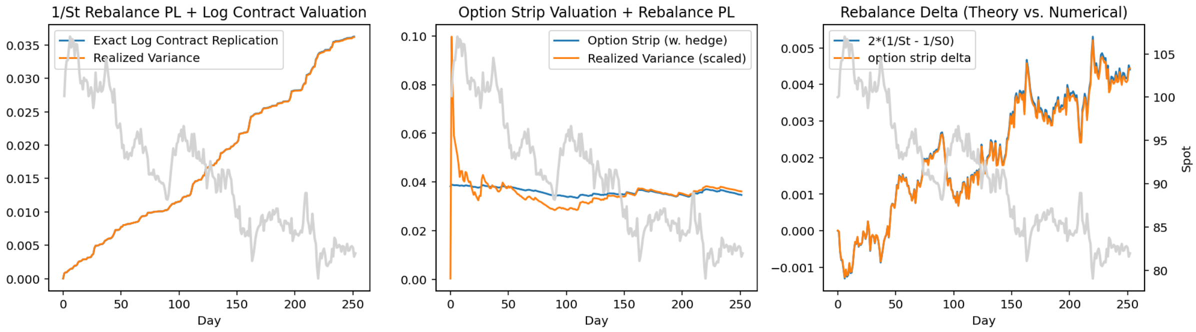

Figure 1: Simulation output for variance swap replication and log-contract decomposition.

A code copy is attached below.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.stats import norm


S0 = 100.0
T = 12.0 / 12.0
N = round(252 * T)
dt = 1.0 / 252.0
r = 0.0
sigma = 0.20

np.random.seed(88)


def bs_price(spot, strike, tau, rate, vol, opt_type):
    if tau <= 1e-12:
        return max(spot - strike, 0.0) if opt_type == "call" else max(strike - spot, 0.0)
    d1 = (np.log(spot / strike) + (rate + 0.5 * vol**2) * tau) / (vol * np.sqrt(tau))
    d2 = d1 - vol * np.sqrt(tau)
    if opt_type == "call":
        return spot * norm.cdf(d1) - strike * np.exp(-rate * tau) * norm.cdf(d2)
    return strike * np.exp(-rate * tau) * norm.cdf(-d2) - spot * norm.cdf(-d1)


def bs_delta(spot, strike, tau, rate, vol, opt_type):
    if tau <= 1e-12:
        if opt_type == "call":
            return 1.0 if spot > strike else (0.0 if spot < strike else 0.5)
        return 0.0 if spot > strike else (-1.0 if spot < strike else -0.5)
    d1 = (np.log(spot / strike) + (rate + 0.5 * vol**2) * tau) / (vol * np.sqrt(tau))
    return norm.cdf(d1) if opt_type == "call" else norm.cdf(d1) - 1.0


def simulate_spot_path(spot0, steps, step_dt, vol):
    log_returns = (-0.5 * vol**2) * step_dt + vol * np.sqrt(step_dt) * np.random.normal(0.0,
        1.0, steps)
```

```python
    spot = spot0 * np.exp(np.cumsum(log_returns))
    return np.concatenate(([spot0], spot))


def build_otm_strip(spot0, strike_min=50, strike_max=200.0, strike_step=1):
    put_strikes = np.arange(strike_min, spot0, strike_step)
    call_strikes = np.arange(spot0 + strike_step, strike_max + strike_step, strike_step)
    put_weights = 2.0 * strike_step / put_strikes**2
    call_weights = 2.0 * strike_step / call_strikes**2
    return put_strikes, put_weights, call_strikes, call_weights


def strip_value_and_delta(spot, tau, put_strikes, put_weights, call_strikes, call_weights):
    put_values = np.array([bs_price(spot, k, tau, r, sigma, "put") for k in put_strikes])
    call_values = np.array([bs_price(spot, k, tau, r, sigma, "call") for k in call_strikes])
    put_deltas = np.array([bs_delta(spot, k, tau, r, sigma, "put") for k in put_strikes])
    call_deltas = np.array([bs_delta(spot, k, tau, r, sigma, "call") for k in call_strikes])
    value = np.sum(put_weights * put_values) + np.sum(call_weights * call_values)
    delta = np.sum(put_weights * put_deltas) + np.sum(call_weights * call_deltas)
    return value, delta

def exact_log_delta(spot):
    return -2.0 / spot


S = simulate_spot_path(S0, N, dt, sigma)
dS = np.diff(S)
realized_var = np.cumsum((dS / S[:-1]) ** 2)

put_strikes, put_weights, call_strikes, call_weights = build_otm_strip(S0)

taus = np.maximum(T - np.arange(N + 1) * dt, 0.0)
log_values = np.zeros(N + 1)
strip_values = np.zeros(N + 1)
linear_values = np.zeros(N + 1)
log_deltas = np.zeros(N + 1)

strip_deltas = np.zeros(N + 1)
linear_deltas = np.full(N + 1, -2.0 / S0)

for i in range(N + 1):
    strip_values[i], strip_deltas[i] = strip_value_and_delta(
        S[i], taus[i], put_strikes, put_weights, call_strikes, call_weights
    )
    linear_values[i] = -(2.0 / S0) * (S[i] - S0)
    log_deltas[i] = exact_log_delta(S[i])


hedge_exact = -log_deltas


### Part1.
dynamic_pnl_exact = np.cumsum(hedge_exact[:-1] * dS)
log_payoff_exact = -2.0 * np.log(S[1:] / S0)
rv_replication_exact = dynamic_pnl_exact + log_payoff_exact

### Part2.
# hedge log contract...
hedge_strip = -strip_deltas   # hedge option strip
hedge_strip_pnl = np.cumsum(hedge_strip[:-1] * dS)
rv_replication_option = hedge_strip_pnl + strip_values[1:]


fig, axes = plt.subplots(1, 3, figsize=(14, 4))

axes[0].plot(rv_replication_exact, label="Exact Log Contract Replication")
axes[0].plot(realized_var, label="Realized Variance")
axes[0].set_title("1/St Rebalance PL + Log Contract Valuation")
axes[0].set_xlabel("Day")
axes[0].legend()
ax0_right = axes[0].twinx()
ax0_right.plot(np.arange(1, N + 1), S[1:], color="lightgrey", linewidth=2)
ax0_right.set_yticks([])
```

```
ax0_right.spines["right"].set_visible(False)

axes[1].plot(rv_replication_option, label="Option Strip (w. hedge)")
axes[1].plot(realized_var*(T/(T-taus[1:])), label="Realized Variance (scaled)")
axes[1].set_title("Option Strip Valuation + Rebalance PL")
axes[1].set_xlabel("Day")
axes[1].legend()
ax1_right = axes[1].twinx()
ax1_right.plot(np.arange(1, N + 1), S[1:], color="lightgrey", linewidth=2)
ax1_right.set_yticks([])
ax1_right.spines["right"].set_visible(False)


axes[2].plot(2*(1/S - 1/S[0]), label="2*(1/St - 1/S0)")
axes[2].plot(-strip_deltas, label='option strip delta')
axes[2].set_title('Rebalance Delta (Theory vs. Numerical)')
axes[2].set_xlabel("Day")
axes[2].legend()
ax2_right = axes[2].twinx()
ax2_right.plot(np.arange(N + 1), S, color="lightgrey", linewidth=2)
ax2_right.set_ylabel("Spot")

plt.tight_layout()
plt.show()
```

References

[AEG06] Peter Allen, Stephen Einchcomb, and Nicolas Granger. Variance swaps. *London: JP Morgan (November)*, 2006.