Realized measures estimation of volatility

Eduardo Rossi

Università di Pavia

Econometria finanziaria
2011

Outline

Stylized facts:

- Actual realizations of financial returns volatility are not directly observable.
- Financial volatility changes through time.
- Both ex-ante and ex-post volatility measures are in common use.

- *Univariate risky log-price process*: $p(t)$ defined on $(\Omega, \mathcal{F}, P)$. The price process evolves in continuous time over the interval $[0, T]$, T finite integer.
- *Natural filtration*: $(\mathcal{F}_t)_{t \in [0,T]} \subseteq \mathcal{F}$. $\mathcal{F}_t$ contains the full history (up to time t) of the realized values of the asset price and other relevant (possibly latent) state variables.
- *Information set* generated by the asset price history alone: $(F_t)_{t \in [0,T]} \subseteq F \equiv F_T$, by definition $F_t \subseteq \mathcal{F}_t$.

The continuously compounded return over the time interval $[t-h, t]$ is

$$r(t,h) = p(t) - p(t-h) \qquad 0 \leq h \leq t \leq T$$

$$r(t) \equiv r(t,t) = p(t) - p(0)$$

$$\begin{aligned} r(t,h) &= p(t) - p(0) + p(0) - p(t-h) \\ &= r(t) - (p(t-h) - p(0)) \\ &= r(t) - r(t-h) \end{aligned}$$

Maintained assumption:

- the asset price process is almost surely strictly positive and finite, so that $p(t)$ and $r(t)$ are well defined over $[0, T]$ a.s.
- $r(t)$ has only countably many jumps points over $[0, T]$.

Càdlàg version of the process for which

$$r(t-) \equiv \lim_{\tau \to t, \tau < t} r(\tau)$$
$$r(t+) \equiv \lim_{\tau \to t, \tau > t} r(\tau)$$

$$r(t) = r(t+) \qquad \text{a.s.}$$

The *jumps* in the cumulative price and return process are

$$\Delta r(t) \equiv r(t) - r(t-) \qquad 0 \leq t \leq T$$

At continuity points for $r(t)$: $\Delta r(t) = 0$.

A jump occurrence is unusual in the sense that we generically have:

$$Pr(\Delta r(t) \neq 0) = 0 \qquad t \in [0, T]$$

This does not imply that jumps necessarily are rare. There is the possibility of a *(countably) infinite number* of jumps over any discrete interval - a phenomenon referred to as an explosion.

Regular Processes: Jump processes that do not explode.
For regular processes, the anticipated jump frequency is conveniently characterized by the instantaneous jump intensity, i.e., the probability of a jump over the next instant of time, and expressed in units that reflect the expected (and finite) number of jumps per unit time interval.
Regular Processes: Jump processes that do not explode.

In a frictionless word, no-arbitrage opportunities and finite expected returns imply log-price process must constitute a semi-martingale (Black, 1991).

The unique canonical return decomposition (Protter, 1992).
Any *arbitrage-free logarithmic price* process (subject to regularity conditions) may be uniquely represented as

$$r(t) \equiv p(t) - p(0) = \mu(t) + M(t) = \mu(t) + M^C(t) + M^J(t)$$

instantaneous return decomposed into an expected return component and a martingale innovation:

- $\mu(t)$ is a predictable and finite variation process
- $M(t)$ local martingale
- $M^C(t)$ a continuous sample path, infinite variation local martingale component.
- $M^J(t)$ a compensated jump martingale.

Assumption: normalized initial conditions

$$\mu(0) \equiv M(0) \equiv M^C(0) \equiv M^J(0) \equiv 0$$

This provides a unique decomposition of the instantaneous return into an expected return component and a (martingale) innovation.

Over discrete intervals,

$$\begin{aligned} r(t,h) &= r(t) - r(t-h) \\ &= \mu(t) - \mu(t-h) + M(t) - M(t-h) \\ &= \mu(t,h) + M(t,h) \end{aligned}$$

where

$$\mu(t,h) \equiv \mu(t) - \mu(t-h) \qquad 0 < h \leq t \leq T.$$

is the expected returns over $[t-h, t]$.

$$M(t,h) = M(t) - M(h)$$

Expected returns

$$m(t,h) \equiv E[r(t,h)|\mathcal{F}_{t-h}] = E[\mu(t,h)|\mathcal{F}_{t-h}] \qquad 0 < h \leq t \leq T$$

The return innovation takes the form

$$r(t,h) - m(t,h) = (\mu(t,h) - m(t,h)) + M(t,h)$$

- The expected return process, even though it is (locally) predictable, may evolve stochastically over the $[t-h, t]$ interval.
- If $\mu(t,h)$ is measurable with respect to $\mathcal{F}_{t-h}$, and thus known at time $t-h$, then the discrete time return innovation reduces to $M(t,h) \equiv M(t) - M(t-h)$.
- Although the discrete-time return innovation incorporates two distinct terms, the martingale component, $M(t,h)$, is generally the dominant contributor to the return variation over short intervals, i.e., for h small.

Let's decompose the expected return process into a purely continuous, predictable finite variation part, $\mu^c(t)$, and a purely predictable jump part, $\mu^j(t)$.

Because the continuous component, $\mu^c(t)$, is of finite variation *it is locally an order of magnitude smaller* than the corresponding contribution from the continuous component of the innovation term, $M^c(t)$.

The reason is - loosely speaking - that an asset earning, say, a positive expected return over the risk-free rate must have innovations that are an order of magnitude larger than the expected return over infinitesimal intervals. Otherwise, a sustained long position (infinitely many periods over any interval) in the risky asset will tend to be perfectly diversified due to a Law of Large Numbers, as the martingale part is uncorrelated. Thus, the risk-return relation becomes unbalanced.

The presence of a non-trivial $M^j(t)$ component may similarly serve to eliminate arbitrage and retain a balanced risk-return trade-off relationship.

Analogous considerations apply to the jump component for the expected return process, $\mu^j(t)$, if this factor is present.

There cannot be a predictable jump in the mean - i.e., a perfectly anticipated jump in terms of both time and size - unless it is accompanied by large jump innovation risk as well, so that $Pr\{M(t) \neq 0\} > 0$. Again - intuitively - if there was a known, say, positive jump then this induces arbitrage (by going long the asset) unless there is offsetting (jump) innovation risk.

Stochastic Volatility Jump Diffusion with Non-Zero Mean Jumps:

$$\mathrm{d}p(t) = \mu + \beta\sigma^2(t) + \sigma(t)\mathrm{d}W(t) + \kappa(t)\mathrm{d}q(t) \quad 0 \leq t \leq T$$

Assumptions:

- $\sigma(t)$ is a strictly positive continuous sample path process (a.s.)
- $W(t)$ denotes a standard Bm,
- $q(t)$ is a pure jump process with

$$q(t) = \begin{cases} 1 & \text{jump in } t \\ 0 & \text{otherwise} \end{cases}$$

- $\kappa(t)$ refers to the size of the corresponding jumps.
- The jump size distribution : $E[\kappa(t)] = \mu_\kappa$ and $Var[\kappa(t)] = \sigma_k^2$
- The jump intensity is assumed constant (and finite) at a rate λ per unit time.

The return components:

$$
\begin{aligned}
\mu(t) &= \mu^c(t) = \mu \cdot t + \beta \int_0^t \sigma^2(s)\mathrm{d}s + \lambda \cdot \mu_\kappa \cdot t \\
M^c(t) &= \int_0^t \sigma(s)\mathrm{d}W(s) \\
M^j(t) &= \sum_{0 \le s \le t} \kappa(s)q(s) - \lambda \cdot \mu_\kappa \cdot t
\end{aligned}
$$

- The concept of an instantaneous return employed in continuous-time models (SDE form) is pure short-hand notation that is formally defined only through the corresponding integral representation.
- Real-time price data are not available at every instant and, due to pertinent microstructure features, prices are invariably constrained to lie on a discrete grid, both in the price and time dimension.
- There is no real-world counterpart to the notion of a continuous sample path martingale with infinite variation over arbitrarily small time intervals (say, less than a second).
- It is only feasible to measure return (and volatility) realizations over discrete time intervals.
- Sensible measures can only be constructed over much longer horizons than given by the minimal interval length for which consecutive trade prices or quotes are recorded.

Volatility seeks to capture the strength of the (unexpected) return variation over a given period of time.
Two distinct features importantly differentiate the construction of all (reasonable) volatility measures.

1. Forecasts of future return volatility. The focus is on ex-ante expected volatility. Search for a model that may be used to map the current information set into a volatility forecast
2. Given a set of actual return observations, the emphasis is on ex-post measurement of the volatility. The (ex-post) realized volatility may be computed (or approximated) without reference to any specific model (nonparametric procedure).

Focus on the behavior of $M(t)$ process (the martingale component in the return decomposition).
Prerequisite for observing the $M(t)$ is that we have access to a continuous record of price data.
Such data are simply not available, even for extremely liquid markets.
The presence of microstructure effects (discrete price grids, bid-ask bounce effects, etc.) prevents from ever getting really close to a true continuous sample path realization.
We focus on measures that represent the (average) volatility over a discrete time interval, rather than the instantaneous (point-in-time) volatility.

General notion of volatility based on the *quadratic variation process* for the local martingale component in the unique semi-martingale return decomposition.

Quadratic variation

Let $X(t)$ denote any (special) semi-martingale. The *unique quadratic variation process*, $[X, X]_t, t \in [0, T]$, associated with $X(t)$ is then formally defined by

$$[X, X]_t \equiv X(t)^2 - 2\int_o^2 X(s-)\mathrm{d}X(s),$$

where the stochastic integral of the adapted càglàd process, $X(s-)$, with respect to the càdlàg semi-martingale, $X(s)$, is well-defined.

If M is a locally square integrable martingale, then the associated $(M^2 - [M,M])$ process is a local martingale,

$$E[M(t,h)^2 - ([M,M]_t - [M,M]_{t-h})|\mathcal{F}_{t-h}] = 0 \quad 0 < h \leq t \leq T.$$

The quadratic variation process represents the (cumulative) realized sample path variability of $X(t)$ over the $[0,t]$ time interval.
Let a sequence of possible random partitions of $[0,T]$, τ_m, be given s.t.

$$\tau_m \equiv \{\tau_{m,j}\}_{j\geq 0}, \quad m = 1,2,\ldots$$

where

$$\tau_{m,0} \leq \tau_{m,1} \leq \cdots$$

satisfy a.s. for $m \to \infty$,

$$\tau_{m,0} \to 0; \quad \sup_{j\geq 1} \tau_{m,j} \to T; \quad \sup_{j\geq 1} (\tau_{m,j+1} - \tau_{m,j}) \to 0.$$

Then, for $t \in [0, T]$,

$$\lim_{m\to\infty}\{\sum_{j\geq 1}(X(t\wedge\tau_{m,j}) - X(t\wedge\tau_{m,j-1}))^2\} \to [X, X]_t$$

where $t \wedge \tau \equiv min(t, \tau)$, and the convergence is uniform in probability. Theoretical notion of ex-post return variability:

Notional volatility (general)

or Actual volatility (Barndorff-Nielsen and Shephard (2002)), over $[t-h, t], 0 < h \leq t \leq T$:

$$\upsilon^2(t, h) \equiv [M, M]_t - [M, M]_{t-h} = [M^c, M^c]_t - [M^c, M^c]_{t-h} + \sum_{t-h<s\leq t} \Delta M^2(s).$$

Notional volatility

Under the maintained assumption of no predictable jumps in the return process, and noting that the quadratic variation of any finite variation process, such as $\mu^c(t)$, is zero, we also have

$$\upsilon^2(t,h) \equiv [r,r]_t - [r,r]_{t-h} = [M^c, M^c]_t - [M^c, M^c]_{t-h} + \sum_{t-h<s\leq t} \Delta r^2(s).$$

the notional volatility equals (the increment to) the quadratic variation for the return series.

Ex-post it is possible to approximate the notional volatility arbitrarily well through the accumulation of ever finely sampled high-frequency squared return.
This approach remains consistent independent of the expected return process.

The notional volatility

- captures the sample path variability of the log-price process over the $[t - h, t]$ time interval.
- incorporates the effect of (realized) jumps in the price process: jumps contribute to the realized return variability and forecasts of volatility must account for the potential occurrence of such jumps

Expected volatility

The Expected Volatility over $[t-h, t], 0 < h \leq t \leq T$, is defined by

$$\begin{aligned} \vartheta^2(t,h) &\equiv E[\{r(t,h) - E(\mu(t,h)|\mathcal{F}_{t-h})\}^2|\mathcal{F}_{t-h}] \\ &= E[\{r(t,h) - m(t,h)\}^2|\mathcal{F}_{t-h}] \end{aligned}$$

Instantaneous Volatility

$$\sigma_t^2 \equiv \lim_{h \to 0} \left[E\left\{ \frac{[M^c, M^c]_t - [M^c, M^c]_{t-h}}{h} \right\} | \mathcal{F}_{t-h} \right]$$

This definition is consistent with the terminology commonly employed in the literature on continuous-time parametric stochastic volatility models.

Although the instantaneous volatility is a natural concept practical volatility measurement invariably takes place over discrete time intervals.

Two approaches for empirically quantifying volatility:

- Procedures based on estimation of parametric models. Alternative parametric models differentiate through:
  1. different assumptions regarding the expected volatility, $\vartheta^2(t, h)$
  2. distinct functional forms
  3. the nature of the variables in the information set, $\mathcal{F}_{t-h}$.
- Nonparametric measurements, that typically quantify the notional volatility, $\upsilon^2(t, h)$, directly.

Both set of procedures differ importantly in terms of the choice of time interval for which the volatility measure applies:

- a discrete interval, $h > 0$
- a point-in-time (instantaneous) measure, obtained as the limiting case for $h \to 0$.

Parametric Volatility Models

Discrete-Time Parametric Models explicitly parameterize the expected volatility, $\vartheta^2(t,h), h > 0$, as a non-trivial function of the time t-h information set, $\mathcal{F}_{t-h}$.

- *ARCH models*: $\mathcal{F}_{t-h}$ depends on past returns and other directly observable variables only.
- *Stochastic Volatility* (SV) models: $\mathcal{F}_{t-h}$ explicitly incorporates past returns as well as latent state variables.

Continuous-Time Volatility Models: explicit parameterization of the instantaneous volatility, σ_2^2 as a (non-trivial) function of the $\mathcal{F}_t$ information set, with additional volatility dynamics possibly introduced through time variation in the process governing jumps in the price path.

Nonparametric Volatility Measurement

Nonparametric measurement utilizes the ex-post returns, or $\mathcal{F}_\tau$, in extracting measures of the notational volatility:

- *ARCH Filters and Smoothers* rely on continuous sample paths, or $M^j \equiv 0$, in measuring the instantaneous volatility.

  Filters: information up to $\tau = t$
  Smoothers: information up to $\tau > t$.
- *Realized Volatility* measures directly quantify the notional volatility over (non-trivial) fixed-length time intervals.

Set up: No price jumps and frictionless market. The asset's logarithmic price process $p(t)$ must be a semimartingale to rule out arbitrage opportunities:

$$\mathrm{d}p(t) = \mu(t)\mathrm{d}t + \sigma(t)\mathrm{d}W(t) \quad 0 \leq t \leq T$$

- $\mu(t)$ and $\sigma(t)$ are predictable processes,
- $\mu(t)$ is of finite variation,
- $\sigma(t)$ is strictly positive and square integrable, i.e., $E\left[\int_0^t \sigma^2(s)\mathrm{d}s\right] < \infty$
- the processes $\mu(t)$ and $\sigma(t)$ signify the instantaneous conditional mean and volatility of the return.

The continuously compounded return over the time interval from $t-h$ to t,

$$r(t,h) = p(t) - p(t-h) = \int_{t-h}^{t} \mu(s)\mathrm{d}s + \int_{t-h}^{t} \sigma(s)\mathrm{d}W(s)$$

and its quadratic variation $QV(t,h)$ is

$$QV(t,h) = \int_{t-h}^{t} \sigma^2(s)\mathrm{d}s$$

innovations to the mean component $\mu(t)$ do not affect the sample path variation of the return. Intuitively, this is because the mean term, $\mu(t)\mathrm{d}t$, is of lower order in terms of second order properties than the diffusive innovations, $\sigma(t)\mathrm{d}W(t)$.
When cumulated across many high-frequency returns over a short time interval of length h they can effectively be neglected.

- The diffusive sample path variation over $[t-h, t]$ is also known as the integrated variance $IV(t,h)$,
- In this setting, the quadratic and integrated variation coincide. This is however no longer true for more general return process like, e.g., the stochastic volatility jump-diffusion model.

Absent microstructure noise and measurement error, the return quadratic variation can be approximated arbitrarily well by the corresponding cumulative squared return process. Consider a partition

$$\{t-h+\frac{j}{n}, j=1,\ldots,n\cdot h\}$$

$$RV(t,h;n)=\sum_{j=1}^{n\cdot h} r\left(t-h+\frac{j}{n};\frac{1}{n}\right)^2$$

Semimartingale theory ensures that the realized volatility measure converges in probability to the return quadratic variation QV when the sampling frequency n increases:

$$RV(t,h;n) \to QV(t,h) \quad \text{as } n \to \infty.$$

Barndorff-Nielsen and Shephard (2002) showed that for $n \to \infty$:

$$\sqrt{n \cdot h}\left(\frac{RV(t,h;n) - QV(t,h)}{\sqrt{2IQ(t,k)}}\right) \xrightarrow{d} N(0,1)$$

where

$$IQ(t,h) = \int_{t-h}^{t} \sigma^4(s)\mathrm{d}s$$

is the *integrated quarticity*, with $IQ(t,h)$ independent from the limiting Gaussian distribution.

To conduct ex-post inference regarding the actual realized return variation over a given period a consistent estimator of $IQ(t,h)$ is required. $IQ(t,h)$, is unobserved and is likely to display large period-to-period variation.

Barndorff-Nielsen and Shephard (2002) proposed estimators for any integrated power of the diffusive coefficient. The *realized power variation* of order p, $V(p;t,h;n)$, is the (scaled) cumulative sum of the absolute p-th power of the high-frequency returns

$$V(p;t,h;n) \equiv n^{p/2-1}\mu_p^{-1}\sum_{j=1}^{n\cdot h}\left|r\left(t-h+\frac{j}{n};\frac{1}{n}\right)\right|^p$$

where $\mu_p = E|z|^p$ with $z \sim N(0,1)$.

It converges, as $n \to \infty$, to the corresponding power variation of order p, $V(p;t,h)$:

$$V(p;t,h;n) \xrightarrow{p} \int_{t-h}^{t}\sigma^p(s)\mathrm{d}s \equiv V(p;t,h).$$

$V(4; t, h; n)$ is a natural choice as a consistent estimator for the integrated quarticity $IQ(t, k)$. This conclusion is heavily dependent on the absence of jumps in the price process.

Relationship between quadratic variation or integrated variance along with its associated empirical measure, realized volatility, and the conditional return variance.

- In the case of constant drift and volatility coefficients, the conditional (and unconditional) return variance equals the quadratic variation of the log price process.
- when volatility is stochastic we must distinguish clearly between the conditional variance, representing the (ex-ante) expected size of future squared return innovations over a certain period, and the quadratic variation, reflecting the actual (ex-post) realization of return variation, over the corresponding horizon.

Under ideal conditions, the RV captures the actual realizations of return volatility.
One may construct well calibrated forecasts (conditional expectations) of return volatility from a time series of past realized volatilities.

If the instantaneous return is the continuous-time process and the return, mean, and volatility processes are uncorrelated (i.e., $\mathrm{d}W(t)$ and innovations to $\mu(t)$ and $\sigma(t)$ are mutually independent), then $r(t,h)$ is normally distributed conditional on the cumulative drift $\mu(t,h) = \int_{t-h}^{t} \mu(s)\mathrm{d}s$ and the $QV(t,h)$ (which in this setting equals the integrated variance $IV(t,h)$):

$$r(t,h)|\mu(t,h), IV(t,h)) \sim N\left(\mu(t,h), IV(t,h)\right)$$

the return distribution is mixed Gaussian with the mixture governed by the realizations of the IV (and integrated mean) process.
Extreme realizations (draws) from the integrated variance process render return outliers likely while persistence in the integrated variance process induces volatility clustering.

For short horizons, where the conditional mean is negligible relative to the cumulative absolute return innovations, the integrated variance may be directly related to the conditional variance as,

$$Var[r(t,h)|\mathcal{F}_{t-h}] \approx E[RV(t,h;n)|\mathcal{F}_{t-h}] \approx E[QV(t,h)|\mathcal{F}_{t-h}].$$

- A volatility forecast is an estimate of the conditional return variance (on the far left-hand side), which in turn approximates the expected quadratic variation.
- Since *RV* is approximately unbiased for the corresponding unobserved quadratic variation, the realized volatility measure is the natural benchmark against which to gauge the performance of volatility forecasts.
- The quadratic variation is directly related to the actual return variance and to the expected return variance.

- The *RV* concept is associated with the return variation measured over a discrete time interval rather than with the so-called spot or *instantaneous volatility*.
- In principle, the *RV* measurement can be adapted to spot volatility estimation: as $h \to 0$, $QV(t, h) \to \sigma^2(t)$. *RV* converges to instantaneous volatility when both h and h/n shrink.
- For this to happen, however, h/n must converge at a rate higher than h, so as the interval shrinks we must sample returns at an ever increasing frequency.
- In practice, this is infeasible, because intensive sampling over tiny intervals magnifies the effects of microstructure noise.

A broad class of SV models that allow for the presence of jumps in returns is defined by

$$\mathrm{d}p(t) = \mu(t)\mathrm{d}t + \sigma(t)\mathrm{d}W(t) + \kappa(t)\mathrm{d}q(t) \quad 0 \leq t \leq T$$

- where q is a Poisson process uncorrelated with W and governed by the jump intensity $\lambda(t)$:

$$Pr\{\mathrm{d}q_t = 1\} = \lambda_t \mathrm{d}t$$

  with λ_t positive and finite. This assumption implies that there can only be a finite number of jumps in the price path per time period. Common restriction in the finance literature, though it rules out infinite activity Lévy processes.
- The scaling factor $\kappa(t)$ denotes the magnitude of the jump in the return process if a jump occurs at time t.

In this case, the quadratic return variation process over the interval from $t-h$ to t, $0 \leq h \leq t \leq T$, is the sum of the diffusive integrated variance and the cumulative squared jumps:

$$QV(t,h) = \int_{t-h}^{t} \sigma^2(s)\mathrm{d}s + \sum_{s=t-h}^{t} J^2(s) \equiv IV(t,h) + \sum_{s=t-h}^{t} J^2(s)$$

where $J(t) = \kappa(t)dq(t)$ is non-zero only if there is a jump at time t. The *RV* estimator remains a consistent measure of the total QV in the presence of jumps. However, since the diffusive and jump volatility components appear to have distinctly different persistence properties it is useful both for analytic and predictive purposes to obtain separate estimates of these two factors in the decomposition of the quadratic variation implied by equation.

To this end, the *m*-skip bipower variation, BV, introduced by Barndorff-Nielsen and Shephard (2004b) provides a consistent estimate of the *IV* component,

$$BV(t,h;m,n) = \frac{\pi}{2} \sum_{i=m+1}^{n \cdot h} \left| r\left(t - h + \frac{i \cdot h}{n}; \frac{1}{n}\right) \right| \left| r\left(t - h + \frac{(i-m) \cdot h}{n}; \frac{1}{n}\right) \right|$$

Setting $m = 1$ yields the the *realized bipower variation* $BV(t,k;n) \equiv BV(t,h;1,n)$.
The *BPV* is robust to the presence of jumps and therefore, in combination with RV, it yields a consistent estimate of the cumulative squared jump component. As $n \to \infty$:

$$RV(t,h;n) - BV(t,h;n) \to QV(t,h) - IV(t,h) = \sum_{s=t-h}^{t} J^2(s)$$

Huang and Tauchen (2005) a measure of jumps, the **Relative Jump**:

$$RJ_t = \frac{RV_t - BV_t}{RV_t} \tag{1}$$

which is an indicator of the contribution (if any) of jumps to the total within-day variance of the process. $100 \cdot RJ$ is a direct measure of the percentage contribution of jumps, if any, to total price variance. Under the maintained assumptions of no jumps, then asymptotically RV_t and BV_t is independent of RV_t conditional on the volatility path, and thus RJ_t is asymptotically the ratio of two conditionally independent random variables.

Jump Test Statistics Based on Barndorff-Nielsen and Shephard's (2006) theoretical results, Andersen, Bollerslev, and Diebold (2004) use the time series

$$z_t = \frac{RV_t - BV_t}{\sqrt{(v_{bb} - v_{qq})\frac{1}{n}TP_t}}$$

where

$$\begin{aligned} v_{bb} &= (\mu_1^{-1} - 1) + 2(\mu_1^{-2} - 1) \\ v_{qq} &= \mu_4 - \mu_2^2 \end{aligned}$$

Andersen, Bollerslev, and Diebold (2004) suggest using the jump-robust realized TP_t, *Tri-Power Quarticity statistic*, which is a special case of the multipower variations studied in Barndorff-Nielsen and Shephard (2004a)

$$TP_t \equiv n\mu_{4/3}^{-3}\left(\frac{n}{n-2}\right)\sum_{i=3}^{n}\left|r_{t-1,i-2}\right|^{4/3}\left|r_{t-1,i-1}\right|^{4/3}\left|r_{t-1,i}\right|^{4/3}$$

$$r_{t-1,j} \equiv r\left(t-1+\frac{j}{n};\frac{1}{n}\right)$$

with

$$TP_t \xrightarrow{p} \int_{t-1}^{t} \sigma^4(s)\mathrm{d}s$$

even in the presence of jumps. For each t, $z_t \xrightarrow{d} N(0,1)$ as $n \to \infty$. Thus the sequence $\{z_t\}_{t=1}^{T}$ provides evidence on the daily occurrence of jumps in the price process.

There is a scale normalizing constant n in front of the summation because each absolute return is of order $\sqrt{\Delta t}$, so the product is of order $(\Delta t)^2$, and the summation Δt. n is $1/\Delta t$, which cancels out the summation order, and the whole expression approaches a well defined limit.

Suppose the price is contaminated by MN, $\eta_t \sim i.i.d.(0, \omega^2)$:

$$\widetilde{p}_t = p_t + \eta_t$$

with η independent of p_t.

This noise introduces spurious volatility and a negative serial correlation in $\mathrm{d}\widetilde{p}_t$.

Independent market MN leads to a bias that diverges to infinity.

The market microstructure (MN) dynamics generate a divergence between the observed price and the true or *noiseless* price process, whose quadratic variation is the object of interest.

- With microstructure noise the *RV* is both biased and inconsistent. Zhou (1996), Hansen and Lunde (2006), Bandi and Russell (2006) Bandi and Russell (2008).
- Robust Estimators (e.g. TSRV estimator Zhang Ait-Sahalia Mykalnd (2005)).
- Simple solution is to sample at moderate frequencies, e.g., every $5-$, $10-$, or $30-$ minutes.

Hansen and Lunde (2006) consider the following estimator:

$$RV_{HL}^{\Delta} = \sum_{i=1}^{n} r_{i\Delta,\Delta}^{2} + 2\frac{M}{M-1}\sum_{i=2}^{M} r_{i\Delta,\Delta} r_{(i-1)\Delta,\Delta}.$$

which is robust to the first order serial correlation in returns.

The Realized Range, *RRG*, is a rivaling approach to the estimation of *IV*, based on the aggregation of intra-daily ranges.
Martens and Van Dijk (2007), and Christensen and Podolskij (2007) and Christensen, Podolskij and Vetter (2009).

The *RRG* is appealing when MN prevents from the use of the whole record of high-frequency prices to compute *RV* since it exploits a larger amount of information and it is in principle able to attain a higher precision.

The daily range estimator of Parkinson (1980) developed under the assumption that the log-price follows a scaled Brownian motion:

$$\mathrm{d}p(t) = \sigma \mathrm{d}W(t)$$

the daily range is defined as

$$s_p = \sup_{0 \leq t,s \leq 1} \{p_t - p_s\} .$$

On the $[0, 1]$ interval:

$$E[s_p^r] = \lambda_r \sigma^r .$$

Under the assumption of a fully observed continuous time log-price path, Parkinson's estimator of the daily integrated volatility:

$$RG_P = \frac{1}{\lambda_2} s_p^2 = \frac{s_p^2}{4\log(2)}.$$

where

$$\lambda_r = E[s_W^r] = E[\sup\{W_1 - W_0\}^r]$$

Consider an equidistant partition $0 = t_0 < t_1 < \ldots < t_n = 1$, where $t_i = i/n$, and $\Delta = 1/n$.
The intraday range at sampling times t_{i-1} and t_i ($i = 1, 2, \ldots, n$) is

$$s_{p_{i\Delta,\Delta}} = \sup_{t_{i-1} \leq t, s \leq t_i} \{p_t - p_s\}.$$

The *RRG* estimator for the interval $[0, 1]$ is defined as:

$$RRG^{\Delta} = \frac{1}{\lambda_2} \sum_{i=1}^{n} s^2_{p_{i\Delta,\Delta}}$$

Christensen and Podolskij (2007) show that

$$RRG^{\Delta} \xrightarrow{p} IV.$$

Result obtained for very general continuous time processes, including models with leverage, long-memory, diurnal effects or jumps (in $\sigma(t)$).

RRG^{Δ} converges in law to a mixed normal with σ governing the mixture, i.e.:

$$\sqrt{n}(RRG^{\Delta} - IV) \xrightarrow{d} MN(0, \Lambda IQ)$$

with $IQ = \int_0^1 \sigma(u)^4 \mathrm{d}u$.

For RRG Λ is approximately 0.4 while for RV is 2.

RRG^{Δ} uses all the data, whereas RV is based on high-frequency returns sampled at fixed points in time. It follows that

$$\frac{\sqrt{n}(RRG^{\Delta} - IV)}{\sqrt{\Lambda RRQ^{\Delta}}} \xrightarrow{d} N(0, 1)$$

where $RRQ^{\Delta} = \frac{n}{\lambda_4} \sum_{i=1}^{n} s^4_{p_{i\Delta},\Delta}$ consistently estimates the IQ.

When the inference is based on a finite sample, the intraday high-low statistic will be progressively more downward biased as n gets larger, since the number of prices in each Δ decreases.

The true range is not observed!

The source of bias is λ_2, which is constructed on the presumption that p is fully observed.

- Christensen and Podolskij (2007) assume that $mn + 1$ equally spaced price observations are available.
- n intervals each with m returns.
- The log-price for each time in the interval $(0, 1)$ is $p_{\frac{i-1}{n}+\frac{t}{mn}}$, $i = 1 \ldots, n$ and $t = 0, \ldots, m$.
- The observed range over the i-th interval is:

$$s_{p_{i\Delta,\Delta},m} = \max_{0 \leq s,t \leq m} \left\{ p_{\frac{i-1}{n}+\frac{t}{mn}} - p_{\frac{i-1}{n}+\frac{s}{mn}} \right\}.$$

$$s_{W,m} = \max_{0 \leq s,t \leq m} \left\{ W_{t/m} - W_{s/m} \right\}$$

and

$$\lambda_{r,m} = E[s^r_{W,m}].$$

$\lambda_{r,m}$ is the r-th moment of the range of a sBm over $[0, 1]$ when only m increments of the underlying continuous time process are observed.

Numerical simulation to compute $\lambda_{r,m}$

The *RRG* estimator based on discrete observations is

$$RRG^{\Delta}_m = \frac{1}{\lambda_{2,m}} \sum_{i=1}^{n} s^2_{p_{i\Delta,\Delta},m}.$$

RRG^{Δ}_m is a consistent estimator of IV as $n \to \infty$.

RRG_m^Δ is a consistent estimator of IV as $n \to \infty$. If we assume that the log-price follows the SV process and $m \to c \in \mathbb{N} \cup \infty$:

$$\frac{\sqrt{n}(RRG_m^\Delta - IV)}{\sqrt{\Lambda_m RRQ_m^\Delta}} \stackrel{d}{\to} N(0,1)$$

with $\Lambda_m = \frac{\lambda_{4,m} - \lambda_{2,m}^2}{\lambda_{2,m}^2}$, and

$$RRQ_m^\Delta = \frac{n}{\lambda_{4,m}} \sum_{i=1}^{n} s^4_{p_{i\Delta,\Delta},m}.$$

Vetter (2009): With an i.i.d. MN, the *RRG* estimator of *IV* is

$$RRG^{\Delta}_{m,BC} = \frac{1}{\widetilde{\lambda}_{2,m}} \sum_{i=1}^{n} (s_{\widetilde{p}_{i\Delta,\Delta},m} - 2\widehat{\omega}_N)^2$$

where

$$\widetilde{\lambda}_{r,m} = E\left[\left|\max_{t:\eta_{\frac{t}{m}}=\omega,s:\eta_{\frac{s}{m}}=-\omega}\left(W_{\frac{t}{m}} - W_{\frac{s}{m}}\right)\right|^r\right].$$

$Var[\omega^2]$ can be consistently estimated with

$$\widehat{\omega}^2_N = \frac{RV^N}{2N},$$

$N = nm$ is the total number of log-returns.