Measuring volatility of non-normal returns

John Randal
School of Economics and Finance
Victoria University of Wellington
New Zealand

Stylised facts of returns – I

- Returns are heavy-tailed (non-normal)
- Returns are uncorrelated
- The variation in returns (volatility) evolves through time

Model for returns

We assume the returns on a stock can be described by

$$R_t = \mu_t + \sigma_t \epsilon_t \tag{1}$$

where $R_t = \ln P_t - \ln P_{t-1}$ is the daily return, and we assume

- μ_t is small enough to be ignored
- σ_t is a smooth volatility process
- ϵ_t are independent random variables with zero mean and $E(\epsilon_t^2) = 1$.

What happens when ϵ_t is heavy-tailed?

The traditional measure of volatility is the moving standard deviation (historical volatility)

$$s_t = \sqrt{\frac{1}{2p} \sum_{j=-p}^{p} (R_{t-j} - \bar{R}_t)^2}$$

where $\bar{R}_t = \frac{1}{2p+1} \sum_{j=-p}^{p} R_{t-j}$.

Assuming σ_t is constant over the smoothing window and that ϵ_t are Gaussian, s_t^2 is optimal for σ_t^2.

If it is possible that the ϵ_t are heavy-tailed, we should use a robust estimator. Typical daily return data suggests this is necessary.

Robust scale estimation

Good robust scale estimators are not all that well known. We replicated a simulation study (Lax, 1985) and found that robust estimators based on the whole sample make good gains on simple robust estimators, e.g. IQR, MAD.

Rather than use the standard deviation, we could use a robust estimator over a moving window through the data to estimate scale, giving the volatility estimate V_t.

Due to inherent biases, e.g. $E(\mathrm{IQR}) = 1.3490\sigma$ for normal data, we must correct the estimates.

In theory, the correction factor depends on distribution of ϵ_t and the scale estimator used.

Robust volatility estimation

We introduce a scaling factor by appealing to the identity (1) and the assumption that $E(\epsilon_t^2) = 1$. In general

$$\widehat{\epsilon}_t = \frac{R_t}{V_t}$$

will not have unit variance, since for a general distribution of ϵ_t, $E(V_t) = k\sigma_t$ for some $k > 0$, and a scale estimator V_t.

We rescale V_t using $\widehat{\tau} = \text{var}(\widehat{\epsilon}_t)$, and the robust volatility estimator becomes

$$\widehat{\sigma}_t = \widehat{\tau} V_t \qquad (2)$$

We assume that while the variance may be inefficient over $2p + 1$ observations, it will be reliable over the length of the entire series.

ML estimation for the t-distribution

Based on empirical results on daily returns, we form a view on the distribution of ϵ_t and optimise for this situation.

We select the t-distribution with $\nu = 5$ degrees of freedom.

- has finite variance and kurtosis
- in the observed range $3 \leq \nu \leq 9$
- (local) ML estimates available via the EM algorithm
- estimator of similar form to successful robust estimators of scale

We assume ϵ_t has a scaled t_5-distribution and that σ_t is constant over the window $[t-p, t+p]$ and estimate the scale parameter for the sample $R_{t-p}, \ldots, R_{t+p}$ using maximum likelihood.

This is obtained by iterating the EM algorithm (Dempster, Laird and Rubin, 1977) with

$$\hat{\sigma}^2_{t,i+1} = \frac{1}{2p+1} \sum_{j=-p}^{p} q^{(i)}_{t-j} R^2_{t-j} \tag{3}$$

where

$$q^{(i)}_{t-j} = \frac{\nu+1}{\nu-2} \left[1 + \frac{R^2_{t-j}}{(\nu-2)\hat{\sigma}^2_{t,i}} \right]^{-1}$$

for $j = -p, \ldots, p$.

The volatility estimate is taken to be

$$\widehat{\sigma}_t = \widehat{\tau}\widehat{\sigma}_{t,\infty}.$$

where $\widehat{\tau}$ is the sample variance of the standardised returns.

Properties of $\widehat{\sigma}_t$

Very good when ν is small (i.e. close to 5). No worse than s_t when ν is large.

- performs well for simulated t_ν data ($\nu = 3, 5, 9, \infty$)

- seems to perform well for real data

See Randal, Thomson and Lally (*QF*, 2004).

Stylised facts of returns – II

- Returns are heavy-tailed (non-normal)
- Returns are uncorrelated
- The variation in returns (volatility) evolves through time
- *Returns have a disproportionate frequency of zeroes*

Evidence from New Zealand

- Low prices stocks relative to tick size
- Potentially low trading volume (but not always)

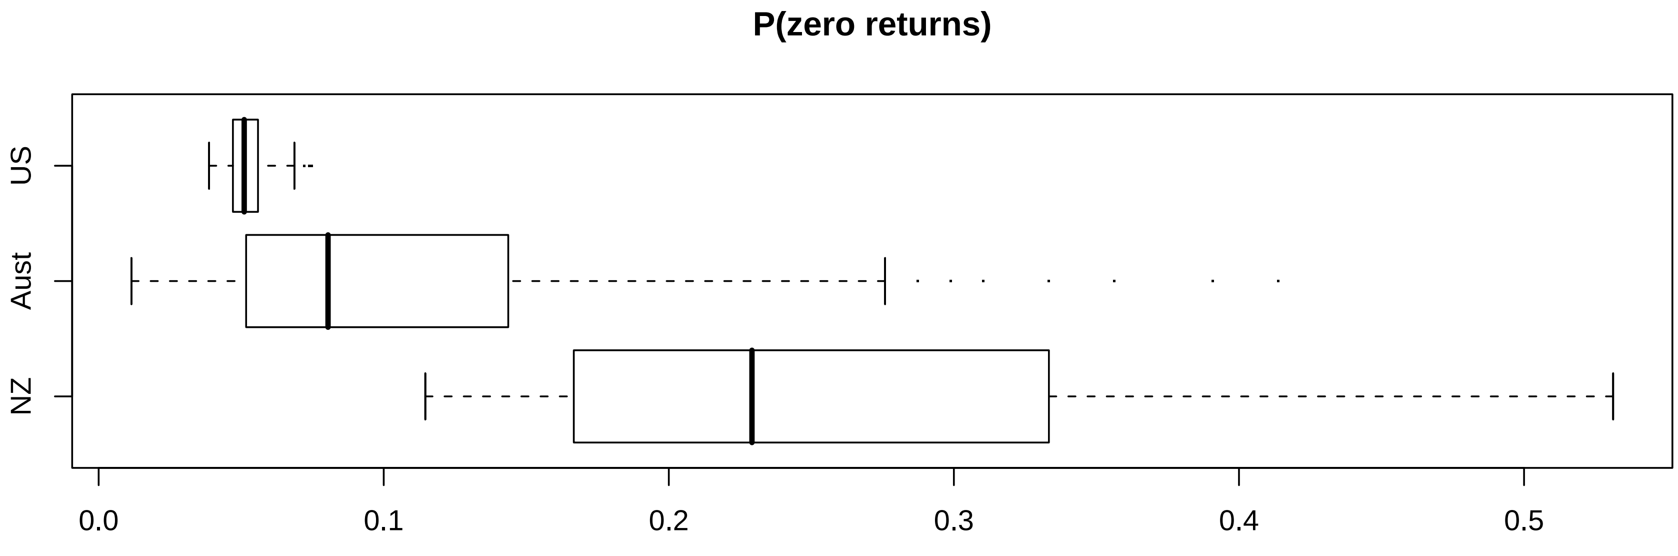

Is rounding important?

Rounded prices are clearly a dominant feature of New Zealand stock prices. Is this important?

We simulate geometric Brownian motion series, round them, and calculate the sample variance of the returns derived from the *rounded prices*. Notably, as series length increases, the quality of the sample variance goes down.

Simulation

We generate (double precision) prices as follows:

$$R_t \sim \text{ iid } N(0.1\Delta, 0.3^2\Delta); t = 1, \ldots, T$$
$$P_0 = 1; P_t = P_0 \exp\left(\sum_{i=1}^{t} R_i\right)$$

where $\Delta = 1$ day $= \frac{1}{250}$ years.

Round prices to the nearest cent (0.01) and form returns R_t^*.

Calculate sample standard deviation for the returns R_t, and the observed returns R_t^*. Repeat 2500 times for fixed T.

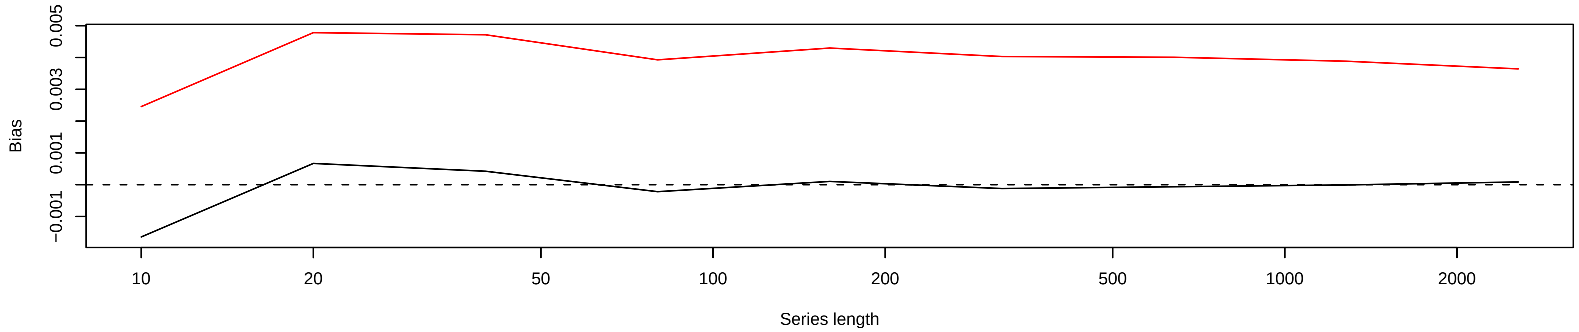
Bias
-0.001
0.001
0.003
0.005
10
20
50
100
200
500
1000
2000
Series length

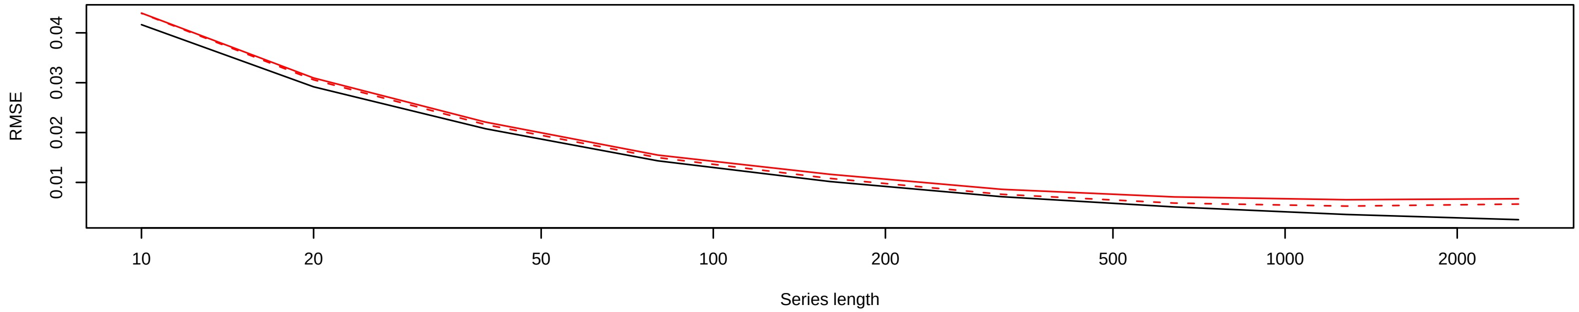
RMSE
0.01
0.02
0.03
0.04
10
20
50
100
200
500
1000
2000
Series length

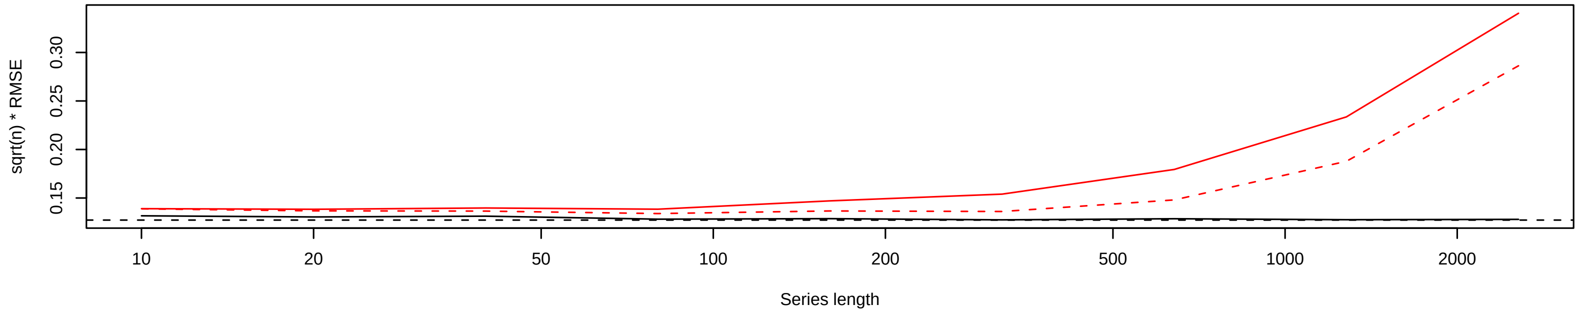
sqrt(n) * RMSE
0.15
0.20
0.25
0.30
10
20
50
100
200
500
1000
2000
Series length

Implications

Recall the correction factor $\hat{\tau} = \text{var}(R_t/V_t)$ and the assumption "that while the variance may be inefficient over [the smoothing window], it will be reliable over the length of the entire series."

In the face of rounded prices, this assumption appears unreasonable, and may worsen as the series length increases.

Developing a correction

Let P_t be the unobserved price.

Let $P_t^* = d\left[\frac{P_t}{d}\right]$ be the observed price, where d is the tick size.

Let $\xi_t = P_t - P_t^*$ be the observation error in the price, with $-\frac{d}{2} \leq \xi_t < \frac{d}{2}$

Let $R_t = \ln P_t - \ln P_{t-\Delta}$ be the unobserved return.

Let $R_t^* = \ln P_t^* - \ln P_{t-\Delta}^*$ be the observed return.

Let $\nu_t = R_t - R_t^*$ be the observation error in the return.

Properties of ν_t - I

$$R_t^* = \ln P_t^* - \ln P_{t-\Delta}^* = \ln\left(1 + \frac{P_t^* - P_{t-\Delta}^*}{P_{t-\Delta}^*}\right)$$

$$= \ln\left(1 + \frac{P_t - \xi_t - (P_{t-\Delta} - \xi_{t-\Delta})}{P_{t-\Delta}^*}\right)$$

$$\approx \frac{P_t - P_{t-\Delta}}{P_{t-\Delta}^*} - \frac{\xi_t - \xi_{t-\Delta}}{P_{t-\Delta}^*}$$

$$\approx R_t - \nu_t$$

where R_t and ν_t are assumed to be independent.

Properties of ν_t - II

While $-\frac{d}{2} \leq \xi_t < \frac{d}{2}$ is bounded, the observation error in the returns is not bounded. Its numerator is, but its demoninator may become very close to zero as the price decreases.

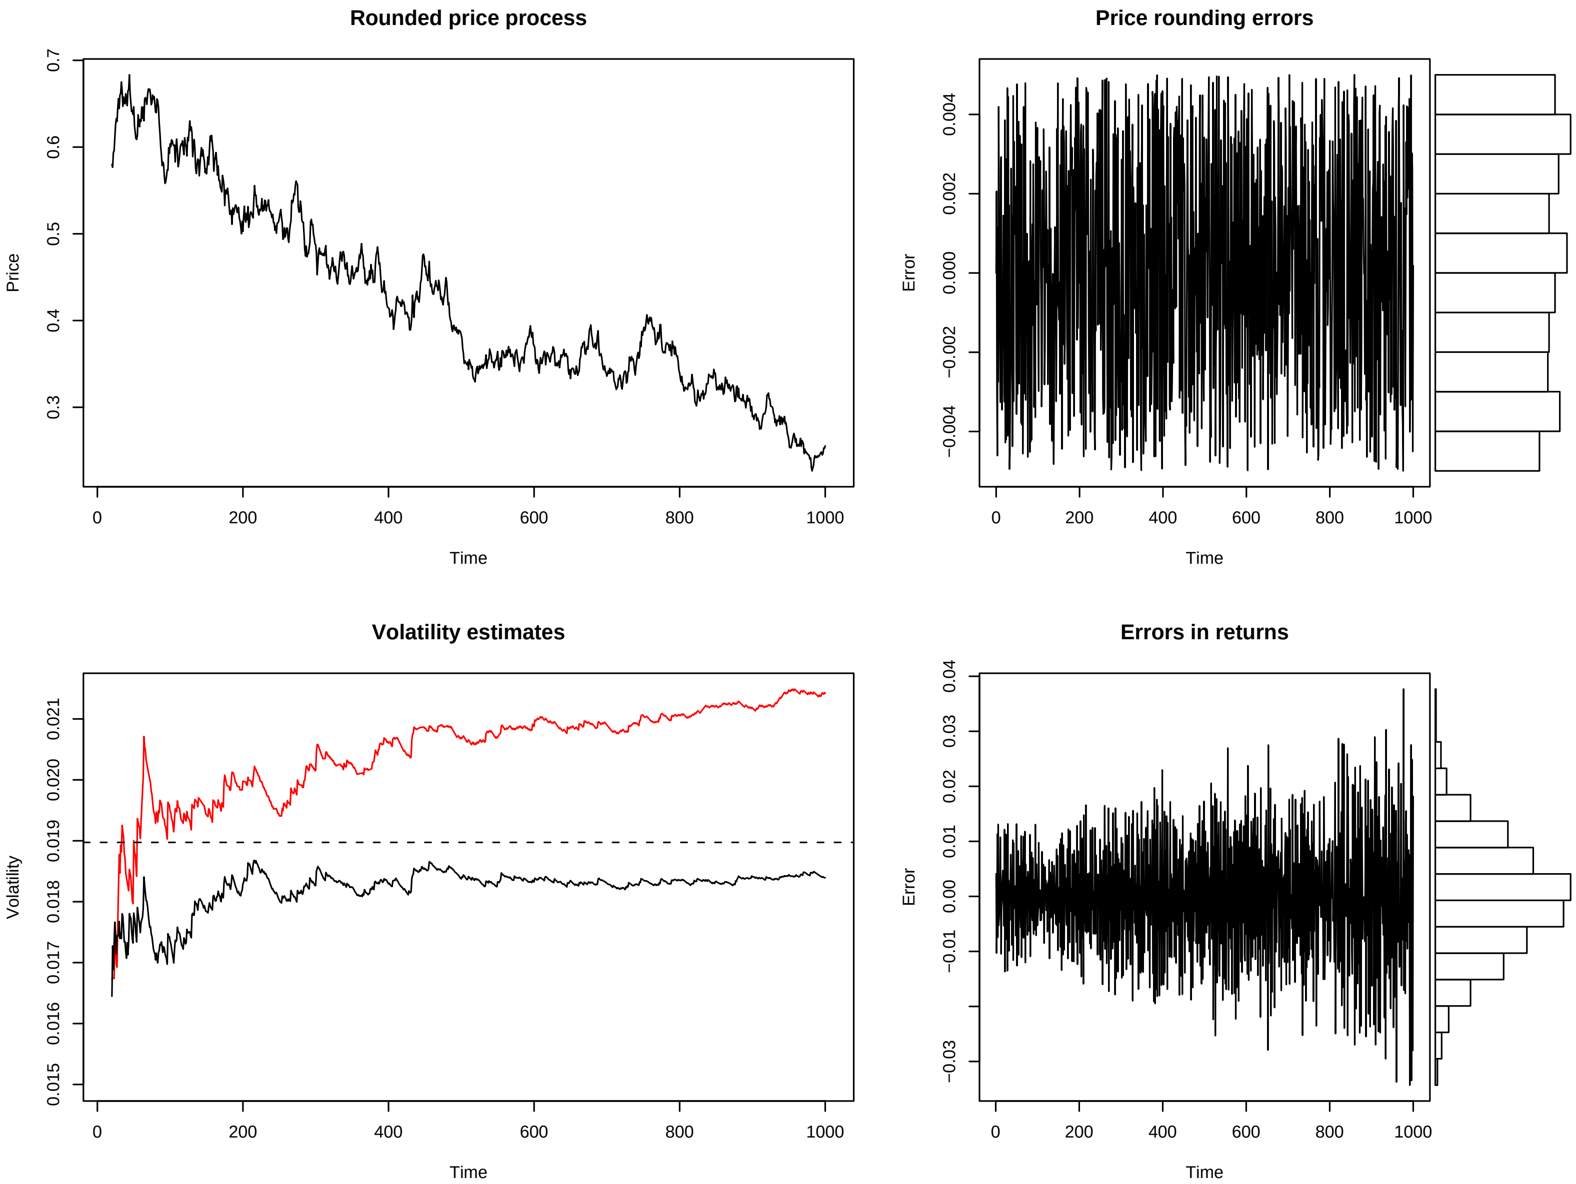
Rounded price process
Price
0.3
0.4
0.5
0.6
0.7
Time
0
200
400
600
800
1000
Price rounding errors
Error
−0.004
−0.002
0.000
0.002
0.004
Time
Volatility estimates
Volatility
0.015
0.016
0.017
0.018
0.019
0.020
0.021
Time
Errors in returns
Error
−0.03
−0.01
0.00
0.01
0.02
0.03
0.04
Time

A volatility correction formula

$$\begin{aligned}
\sigma_*^2 &= \mathsf{var}(R_t^*) = \mathsf{var}(R_t - \nu_t) = \mathsf{var}(R_t) + \mathsf{var}(\nu_t) \\
&= \sigma^2 + \mathsf{var}\left(\frac{\xi_t - \xi_{t-\Delta}}{P_{t-\Delta}^*}\right) \\
&= \sigma^2 + 2\sigma_\xi^2(1-\rho_\xi)\mathsf{E}\left(\frac{1}{P_{t-\Delta}^{*2}}\right)
\end{aligned}$$

$$\sigma^2 = \sigma_*^2 - 2\sigma_\xi^2(1-\rho_\xi)\mathsf{E}\left(\frac{1}{P_{t-\Delta}^{*2}}\right)$$

Evaluating the correction

It seems reasonable to assume $\sigma_\xi^2 = \frac{d^2}{12}$ based on a uniform distribution.

The expectation in the correction term is undefined, however, we investigate using a sample equivalent.

The first-order autocorrelation of ξ_t does not have an obvious theoretical derivation. We consider setting $\rho_\xi = 0$. In addition, we seek a function which is approximately zero when prices are high, and close to one when prices are low. We consider using the proportion of zero returns in the sample to estimate this function.

Estimation of ρ_ξ

We consider short series with initial values P_0 and $T = 125$.

We simulate and record the proportion of zero returns, and the first order sample autocorrelation estimate of the ξ_t series.

We average the proportions, and sample autocorrelations across 2500 series.

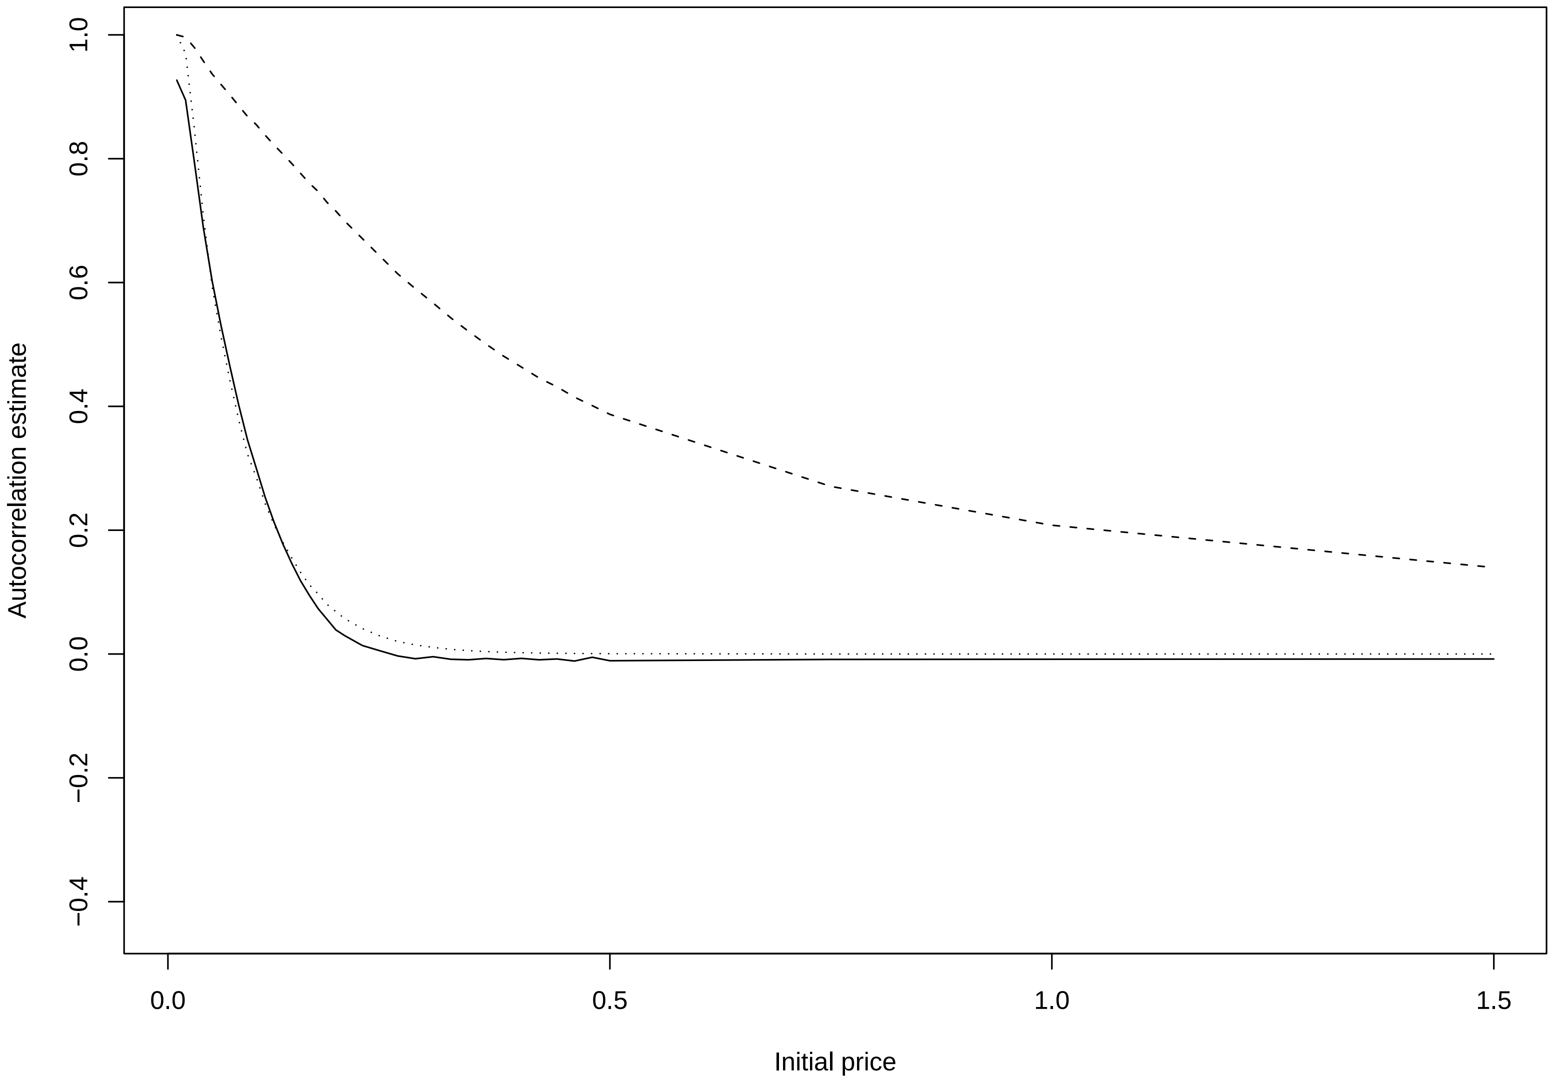

Autocorrelation estimate
1.0
0.8
0.6
0.4
0.2
0.0
−0.2
−0.4
0.0
0.5
1.0
1.5
Initial price

Implementing a correction formula

We recommend an estimator of variance based on rounded prices given by

$$\hat{\sigma}^2 = \frac{1}{T}\sum_{t=1}^{T} R_t^{*2} - \frac{d^2}{6}\left(1 - \left[\frac{\#\{R_t^* = 0\}}{T}\right]^8\right)\frac{1}{T}\sum_{t=1}^{T}\frac{1}{S_{t-\Delta}^{*2}}$$

where prices are rounded to the nearest multiple of d.

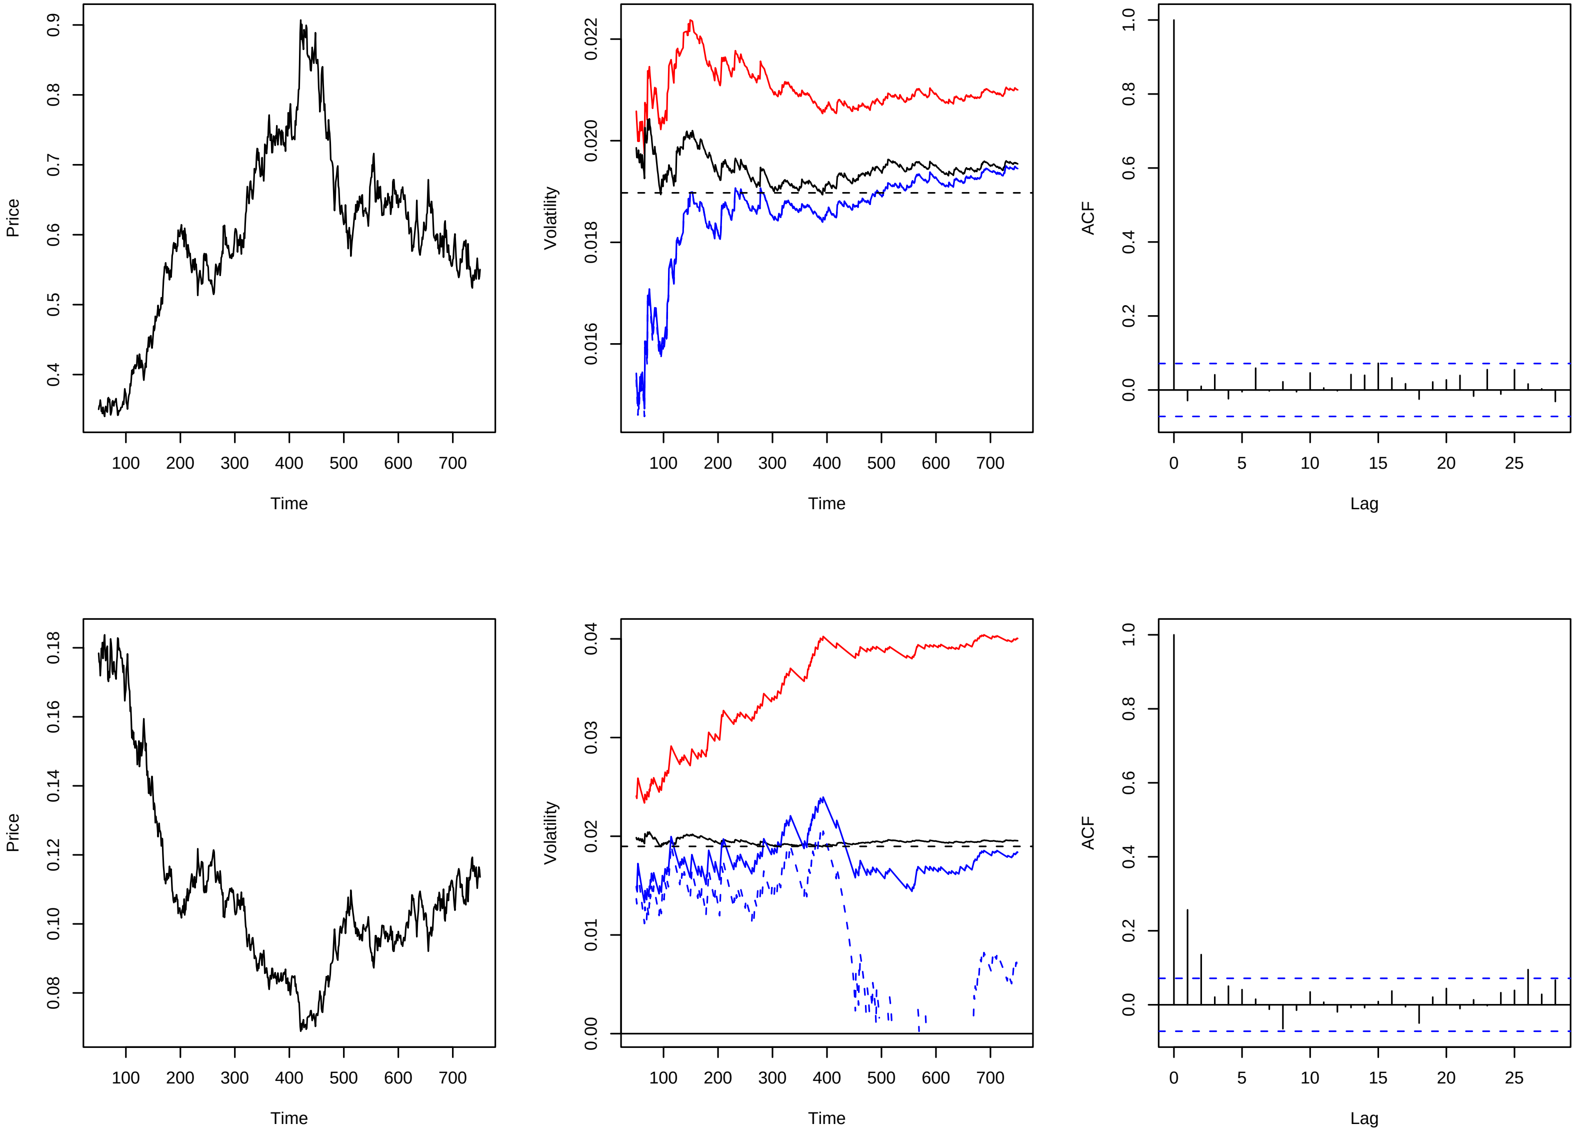

Price
Time
Volatility
Time
ACF
Lag
Price
Time
Volatility
Time
ACF
Lag

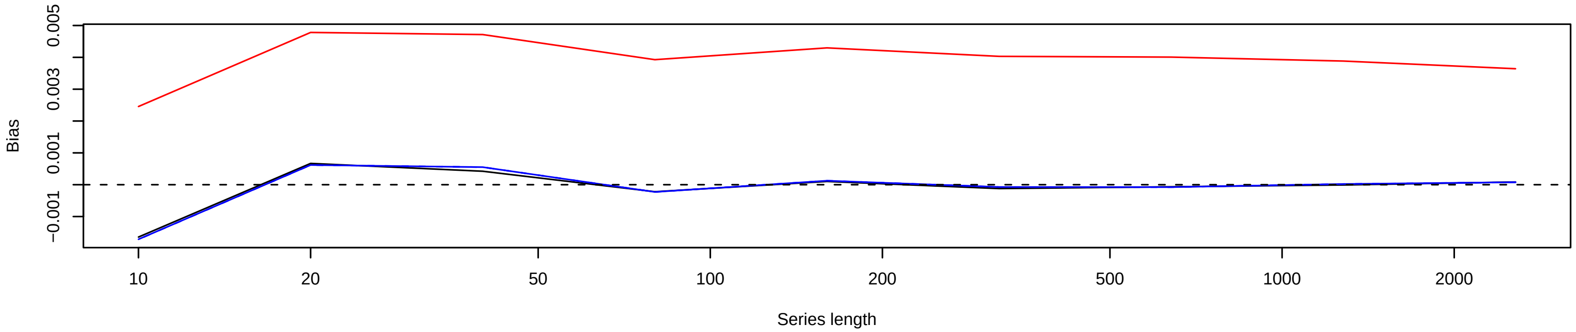
Bias
−0.001
0.001
0.003
0.005
10
20
50
100
200
500
1000
2000
Series length

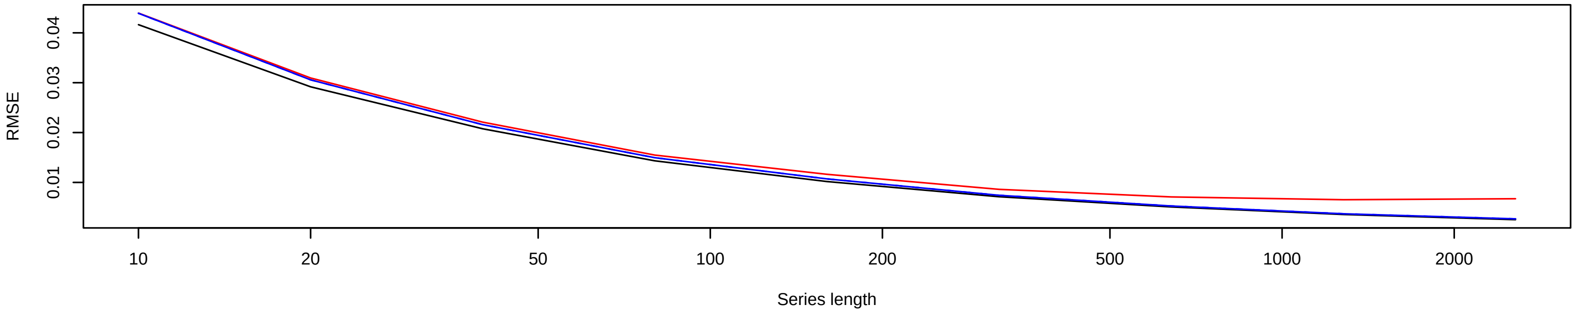
RMSE
0.01
0.02
0.03
0.04
10
20
50
100
200
500
1000
2000
Series length

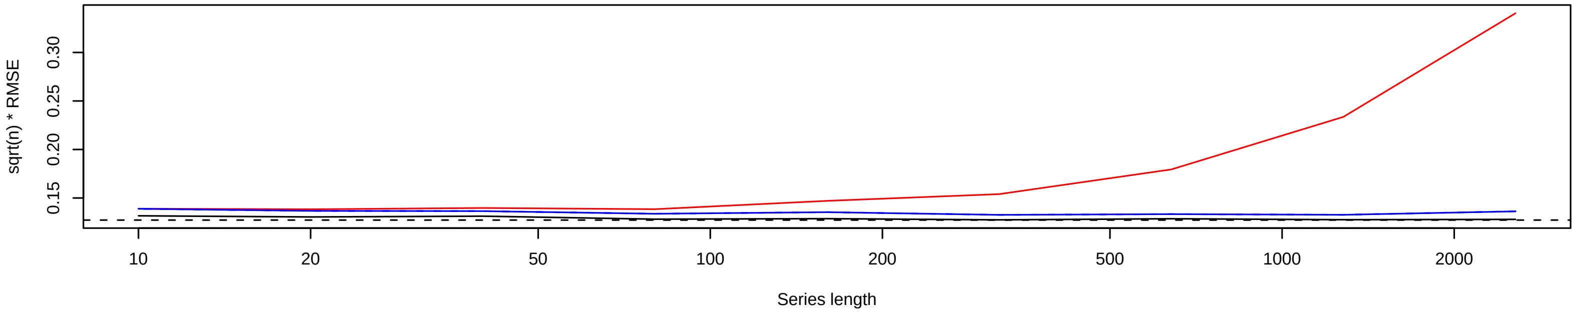
sqrt(n) * RMSE
0.15
0.20
0.25
0.30
10
20
50
100
200
500
1000
2000
Series length

Doing it properly!

We could form the exact likelihood of the observed prices, with

$$f_t(x) = \int_{\log P_{t-1} - \frac{d}{2}}^{\log P_{t-1} + \frac{d}{2}} \frac{1}{\sigma} \phi\left(\frac{x - y - \mu}{\sigma}\right) f_{t-1}(y) dy$$

and

$$l(\mu, \sigma^2; \mathbf{P}) = \int_{\log P_T - \frac{d}{2}}^{\log P_T + \frac{d}{2}} f_T(x) dx$$

and maximise this numerically.

An alternative based on the EM algorithm

Form the exact likelihood based on the complete information set including the (unobserved) prices P_t. This is the standard lognormal likelihood function.

Maximise the expected value of this likelihood, given the data P_t^*, $t = 1, \ldots, T$.

We will need to evaluate

$$E_0(R_t | P_t^*, t = 1, \ldots, T)$$

and

$$E_0\Big((R_t - \widehat{\mu})^2 | P_t^*, t = 1, \ldots, T\Big)$$

These expectations will not be easy to evaluate.

Conclusions

- The empirical correction formula is incredibly ad hoc, and needs theoretical support and/or extensive simulation-based support.

- Nonetheless, this is a very real issue for practitioners in markets like New Zealand's where stock prices are not large relative to tick size.

- There are implications potentially not only for volatility, but other measure estimated from observed (rounded) prices, e.g. CAPM betas.