There’s a hair in my soup!!!

The unacknowledged source of bias in our most common measures of portfolio risk

April 15, 2025

Setting the Context: Parametric vs Nonparametric Methods

- These results pertain to traditional ***parametric methods*** for assessing portfolio risk.
- Nonparametric methods (sampling distributions, hybrid methods) are increasingly popular and have become the preferred approach in some contexts --- Though I would also suggest that parametric methods have been somewhat handicapped in the “horse race” between the two.

Still, if you’ve settled on alternative methods, why should you care?

Why might we care?

- The approach can also be applied to empirical distributions and other hybrid methods (Tsafack and Cataldo, Empirical Economics 2020).
- Also, the approach may can offer an improved alternative metric for validation and cross checking your primary method.
- Now, back to the Feature Presentation…

There’s an egregious error at the heart of traditional measures of portfolio risk!

- The moments of nonlinear transformations of random variables are not equal to the transformation of their moments.
- Many of our most commonly used financial metrics – VaR, Expected Shortfall, Sharpe Ratio, and option values – are nonlinear transformations of sample volatility (which, of course, is a random variable)
- So the expected values we traditionally calculate for all these risk measures are… ***wrong***!!!

It’s one thing to know something is wrong…

- ***But quite another to be able to correct it.***
- This led me on a multi year journey, beginning with a paper published by Georges Tsafack and me in 2020.*
- We showed the distortion in VaR resulting from the naive transformation of moments is substantial and “one-way” – a systematic ***underestimation*** of risk
- However, we were only able to demonstrate and quantify the bias under limited conditions

*Backtesting and estimation error: value-at-risk overviolation rate 2020 Empir Econ DOI 10.1007/s00181-020-01905-4

The quest for practical application takes us well beyond the confines of *VaR*

- What we found is that the main driver of bias is the variability of the volatility parameter: ***The higher the variability of sample vol, the greater the bias.***
- But how do we arrive at a sound estimate of this variability?
- This depends on a number of factors that vary considerably in application, such as the assumed ***distribution of returns***, the ***form of the volatility estimator***, and ***estimator parameterization***.
- How do we tell a coherent story about these factors and weave it into a feasible estimator of the key to it all – the variability of our volatility forecast?

How does the bias principle operate in the specific case of *VaR*

- Here we have the expression for *VaR* as the inverse of the Normal CDF evaluated at the sample estimate of the standard deviation of returns:

$$VaR_{\propto} = \varphi^{-1}(\alpha, 0, s)$$

- This is clearly a nonlinear transformation of a random variable (sample volatility)
- This expression ***does not*** represent the expected value (first moment) of Value-at-Risk

Even under ideal conditions, the resulting bias in VaR is substantial and systematically ***negative***

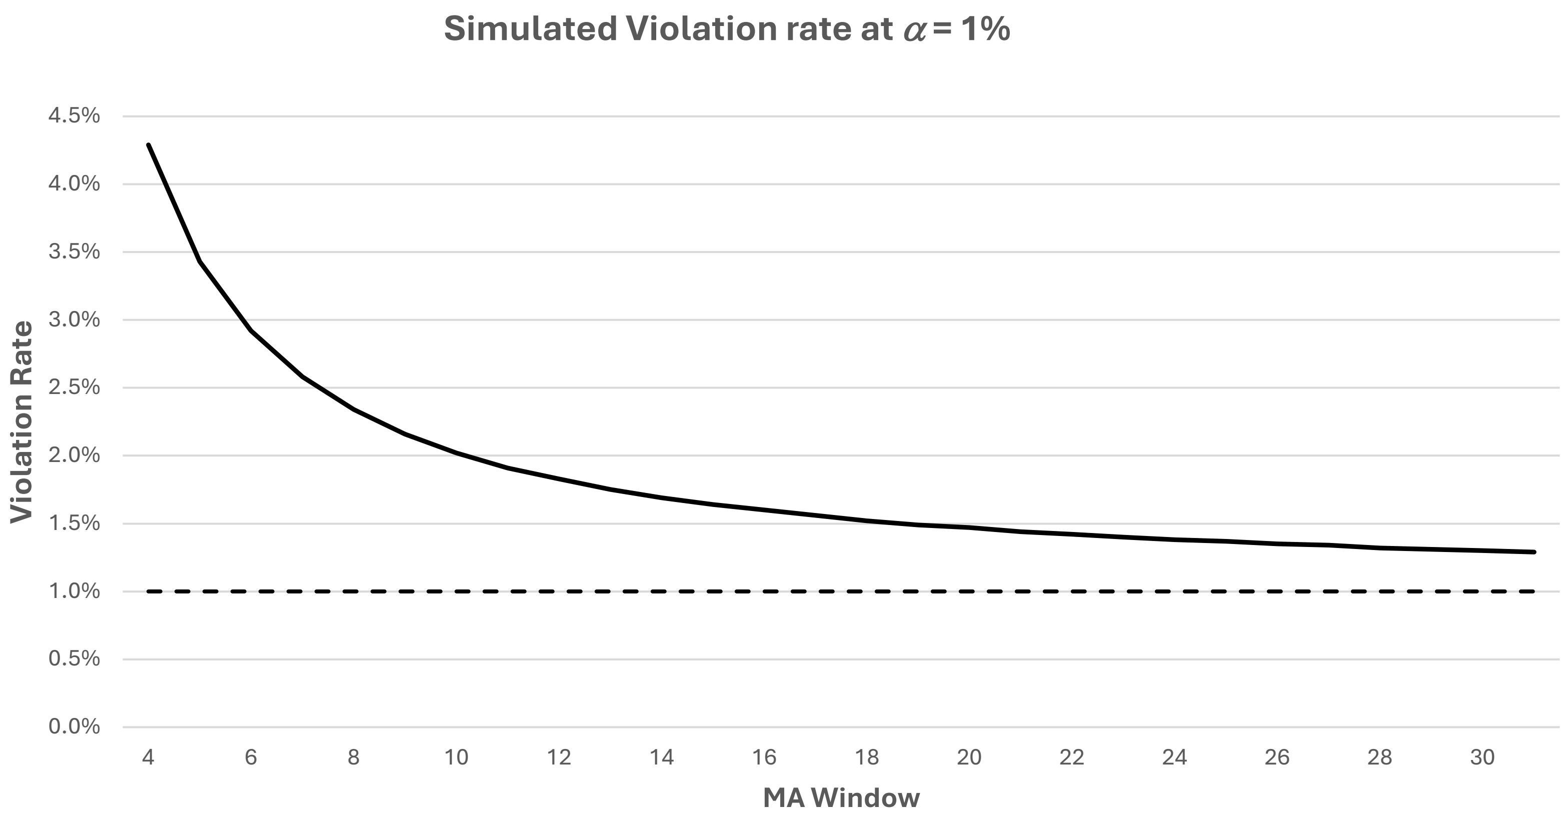

Baseline assumptions for this initial demonstration

- Normal iid returns
- Apply simple MA(n) “forecast” of next period volatility
- This vol “forecast” is unbiased with known distribution of *chi(n)*

Violation rate expressed a result of the joint distribution of sample volatility and returns

$$\Pr(loss > VaR_{\propto}) = \int_{0}^{\infty} chi(k,x) \int_{-\infty}^{VaR_{\propto}*\frac{x}{\mu(k)}} n(y,0,1)\,dy\,dx$$

where

chi(k,x) is the chi distribution with k degrees of freedom

$\mu(k)$ is the mean of the chi(k) distribution

n(y,0,1) is the standard normal distribution

Finding the bias correction factors to conventional *VaR* calculations

Simply solve for the constant "λ" that returns the target violation rate

$$\int_{0}^{\infty} chi(k,x) \cdot \left(\int_{-\infty}^{\lambda \cdot \frac{x}{\mu(k)}} dnorm(y,0,1)\,\mathrm{d}y \right) \mathrm{d}x = .01$$

What do violation rates for "unadjusted" *VaR* look like with Market Data?

MA(10) estimation, EWMA and Garch(1,1) produce comparable results

	Unadjusted Violation Rate	*95% Lower Bound*	*95% Upper Bound*
SPX:w=10	3.02%	*2.38%*	*3.78%*
BBB w=17	2.05%	*1.53%*	*2.69%*
XOI (Oil) w=15	2.55%	*1.96%*	*3.25%*

Challenges to application

We’ve shown that bias exists, it’s always underestimates risk, and it increases with the variance of our volatility forecast.

But…

- The variance of our vol forecast depends on:
  1. The distribution of returns
  2. Our chosen method of estimation (MA(n), EWMA(λ), ARCH, Garch, etc.)
  3. *And* the estimation Parameters.

Which Distribution?

Various schools of thought (diBartolomeo 2007 and others)

1. “Stable” distributions, characterized by infinite variance.
2. Specific, nonnormal distributions with moments reflecting observed behavior of returns.
3. Returns are normal at each instant of time, but with time-varying volatility (Heteroscedastic Normal)
4. Various combinations: Heteroscedastic Non-normal, Composite distributions

My case for pursuing option number three…

- Heteroskedastic Normal is simplest in concept and application
- Strong theoretical case for application to portfolios and indexes (Central Limit Theorem, anyone?)
- Alternatives may be found to work well in specific contexts, but are highly subject to the “curve fitting” problem and unstable estimates of higher moments

Which Vol Forecast Method?

- Are more complex approaches necessarily better?
- ARCH/Garch models are no more than forms of weighted average – no different in kind from simple Moving Average or EWMA
- In nearly all the “horse races” between methods in the literature, MA and EWMA contestants use arbitrarily set lookback or decay factors, whereas the parameterization of more exotic methods is optimized over the sample.
- When the playing field is leveled, the advantage of more exotic methods tend to diminish or disappear.
  - For example, the market data violation rates shown earlier are closely similar using MA, EWMA, and Garch forecasts

Comparisons of Forecast Efficiency

- From "Estimating Volatility" by Stephen Marra
  The Journal of Portfolio Management Quantitative Tools 2023
- MA, EWMA, and Garch(1,1) are in a virtual tie by the MSE Estimation Error Criterion!

MA(3 month):	.0067
EWMA:	.0063
Garch(1,1):	.0063

Trade-offs between Efficiency and Bias

- Sampling error of the estimate decreases with number of observations
- MA, EWMA, Arch, Garch and variations are forms of weighted average of past observations
- Past observations embed a potential bias, since the vol in past periods may differ from the expected vol for the forecast period
- A central supposition of most approaches is ***The relevance of past observations diminishes with distance from the present***
- So while the vol forecast remains, ex ante, unbiased, ***the expected value of bias^2 is positive and increasing in n.***

Finding an estimator for Total Variability

- We express total variability as the Mean Squared Error (MSE)
- MSE can be broken down to the sum of sampling error and bias:

$$MSE = (s^2 + bias^2)$$

- Bias adjustments incorporating s^2 only therefore represents a ***lower bound*** on the VaR overviolation correction
- Before trying to tackle the bias component, how does an “s^2 only” adjustment perform on market data?

Violation Rates Adjusted for Sampling Error Only

	Adjusted for s^2 only	*95% Lower Bound*	*95% Upper Bound*
SPX:w=10	1.91%	*1.40%*	*2.53%*
BBB w=17	1.48%	*1.04%*	*2.04%*
XOI (Oil) w=15	1.55%	*1.10%*	*2.12%*

Extending to Total Variability

- We start with two suppositions already embedded in our choice of estimation method:
    1. The bias contribution to MSE increases with distance in time
    2. The parameterization of the volatility estimator optimizes this trade-off

Example: *MSE* of an MA(*n*) Forecast

- The choice of an ***n*** period lookback implies that ***n*** represents, to the best of the forecaster's knowledge, the best trade-off between efficiency and bias.
- In other words, ***n*** minimizes *MSE* , or $s^2 + bias^2$
- *MSE* minimization implies the marginal reduction in s^2 from an additional observation is approximately equal to the marginal increase in $bias^2$.
- If the estimation window optimizes the tradeoff, and $bias^2$ increases linearly with n, then $bias^2 = \alpha * n$

Solving for bias^2

Supposing $bias^2 = \alpha * n$, minimization of MSE implies

$$\alpha = -d(sampling\ error^2)/dn$$

thus

$$bias^2 = -n * d(sampling\ error^2)/dn$$

Illustration where N=10 represents best trade-off between sampling error and recency

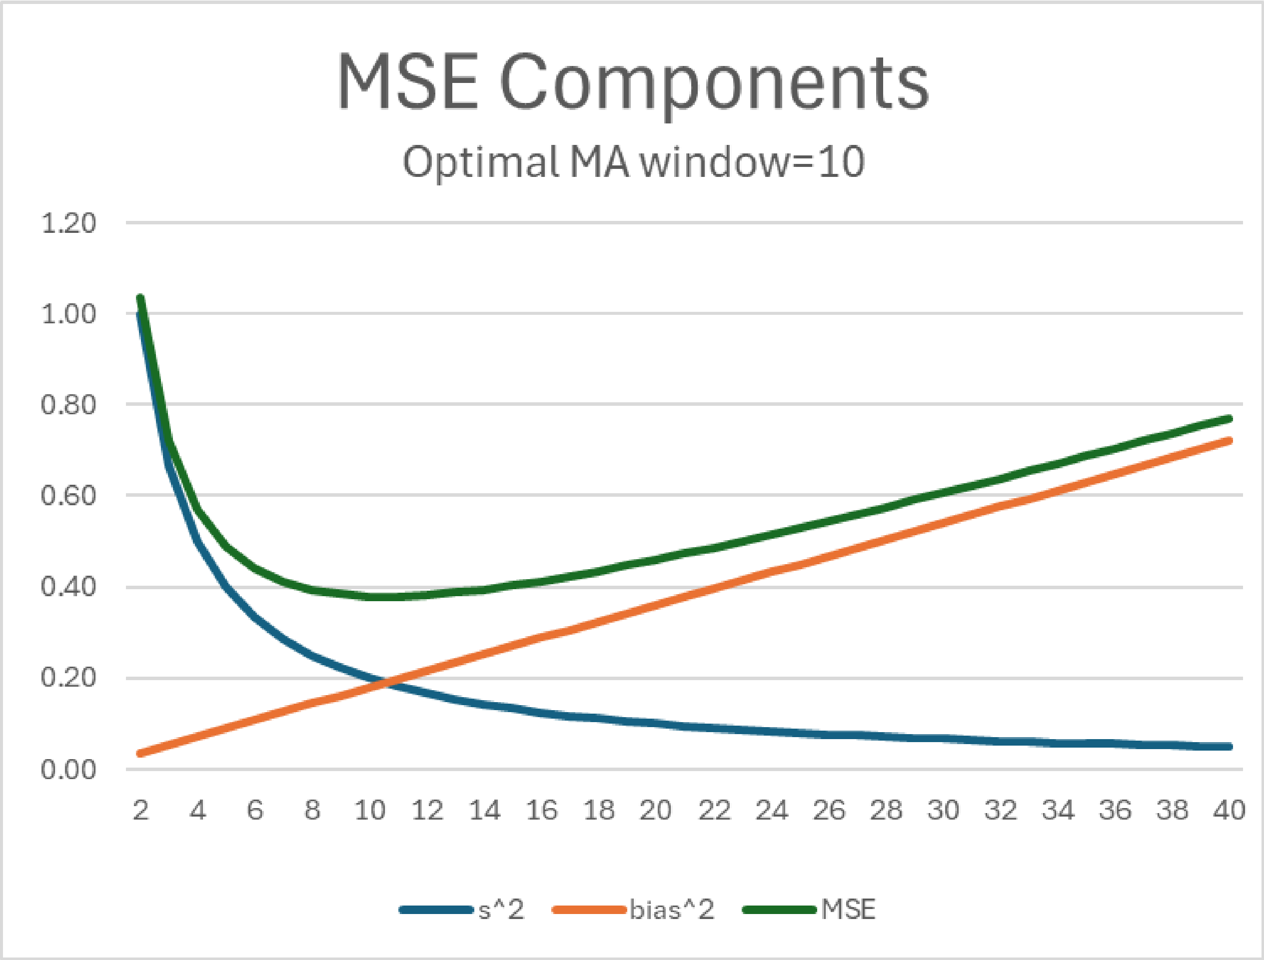

Further steps to calculation of unbiased VaR

- We need to incorporate our new estimate of MSE into the formula
- What’s an appropriate distributional form for sample volatility under conditions of both bias and sampling error?
- We take the chi distribution as a reasonable starting point
- We calibrate degrees of freedom of the chi distribution to match our estimate of MSE.
- Then just solve the integral for λ

Violation Rates

α =.01, Full adjustment for MSE

	Violation Rate	***95% Lower Bound***	***95% Upper Bound***
SPX:w=10	0.83%	*0.51%*	*1.28%*
BBB w=17	0.94%	*0.60%*	*1.41%*
XOI (Oil) w=15	0.96%	*0.61%*	*1.43%*

So how hard is this to do?

- Simple adjustment factor tables can be compiled for different values of MA window λ
- Similar tables can be (less easily) compiled for and EWMA and combinations of Garch parameters
- Implementation is as simple as applying the factor corresponding to the method and parameters used to forecast volatility.
- Preliminary application to Expected Shortfall suggest underestimation effects of similar magnitude.

Beyond *VaR and Expected Shortfall*

- Any financial metrics involving nonlinear transformation of sample estimates or forecasts (volatility, expected returns…)
- The distribution of financial returns, writ large
- What happens when we apply the method over the entire distribution of returns?

Gauss's Revenge?

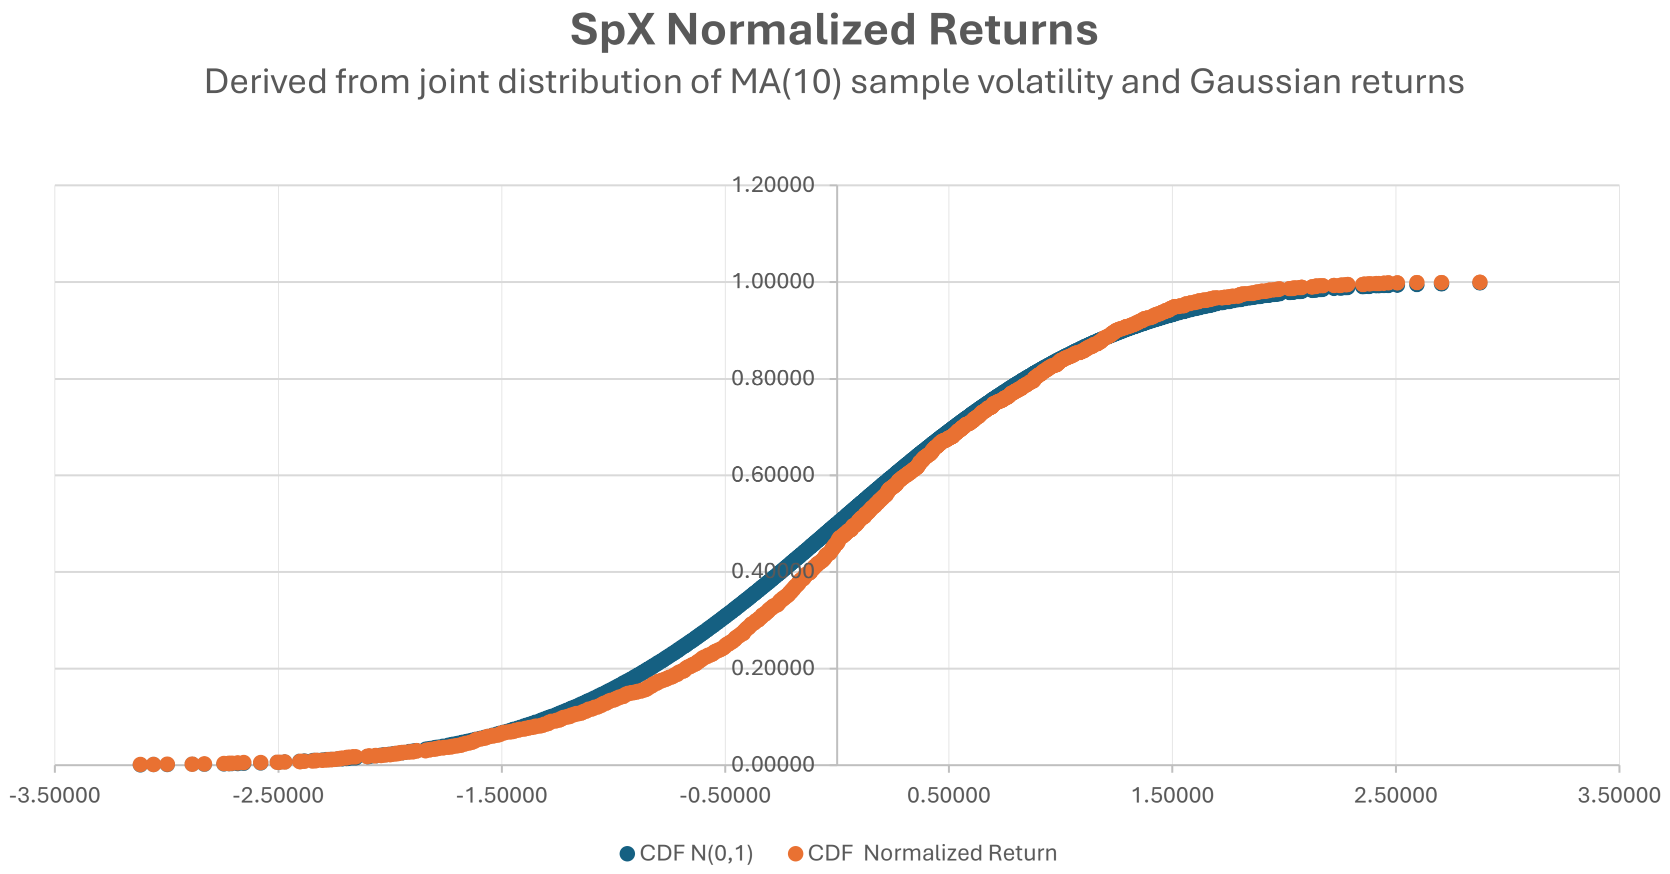

Your thoughts?