Market Risk

MFM Practitioner Module: Quantitiative Risk Management

John Dodson

February 10, 2016

Introduction

This week's material ties together our discussion going back to the beginning of the fall term about risk measures based on the (single-period) loss distribution. Let's recall the taxonomy.

- accounting / capital metric (e.g. mark-to-market loss)
- analysis horizon (e.g. one day, two weeks)
- probability measure: valuation-based, forecast-based, conditional, or equilibrium
- method: analytical, historical simulation, Monte Carlo

Let's also recall the axiomatic discussion from last week about the risk measure itself: ideally law invariant, translation invariant, monotonic, comonotone additive, subadditive, and positive homogeneous.

Loss Operator

Let the net asset value of a portfolio (with fixed holdings) at time τ be $V(\tau)$ and let $\tau_t = t\,(\Delta t)$ where Δt is the natural sampling interval, e.g. business daily. Furthermore let $V_t \triangleq V(\tau_t) \triangleq g(\tau_t, \boldsymbol{Z}_t)$ which represents a mapping of d-dimensional risk factors levels (and time). Finally, let the risk factor innovations be $\boldsymbol{X}_t \triangleq \boldsymbol{Z}_t - \boldsymbol{Z}_{t-1}$.

Loss Operator

Putting all of this together allows us to define the loss in terms of the risk factor innovations,

$$l_{[t]}(\boldsymbol{x}) \triangleq g\left(\tau_t, \boldsymbol{z}_t\right) - g\left(\tau_t + \Delta t, \boldsymbol{z}_t + \boldsymbol{x}\right)$$

so the single-period loss random variable is

$$L_{t+1} = l_{[t]}\left(\boldsymbol{X}_{t+1}\right)$$

Delta-Gamma Approximation

Typically the map g is non-linear, but sometimes it can be well-approximated by a low-order Taylor's expansion about $(\tau_t, \boldsymbol{z}_t)$. The Delta-Gamma approximation is 2$^{\text{nd}}$-order in risk factors and 1$^{\text{st}}$-order in time. We write this as

$$l_{[t]}^{\Delta\Gamma}(\boldsymbol{x}) = -\left(g_\tau\left(\tau_t, \boldsymbol{z}_t\right)\Delta t + \boldsymbol{\delta}\left(\tau_t, \boldsymbol{z}_t\right)' \boldsymbol{x} + \boldsymbol{x}'\Gamma\left(\tau_t, \boldsymbol{z}_t\right)\boldsymbol{x}\right)$$

in terms of

$$\boldsymbol{\delta}(\tau, \boldsymbol{z}) \triangleq \frac{\partial g(\tau, \boldsymbol{z})}{\partial \boldsymbol{z}'} \qquad \Gamma(\tau, \boldsymbol{z}) \triangleq \frac{\partial^2 g(\tau, \boldsymbol{z})}{\partial \boldsymbol{z}' \partial \boldsymbol{z}}$$

and the time decay $g_\tau(\tau, \boldsymbol{z}) = \partial g(\tau, \boldsymbol{z})/\partial\tau$.
Note that the linearity in time is justified by the fact that the components of $\boldsymbol{x}$ have the same order as $\sqrt{\Delta t}$ under a diffusion model for $\boldsymbol{Z}(\tau)$.

Variance-Covariance Method

In the original RiskMetrics model, the authors make the assumption that the risk factor innovations are conditionally normal, $\boldsymbol{X}_{t+1}|\mathcal{F}_t \sim \mathcal{N}_d\left(\boldsymbol{\mu}_{t+1}, \Sigma_{t+1}\right)$, and that the loss operator is affine (effectively $\Gamma = 0$): $l^{\Delta}_{[t]}(\boldsymbol{x}) = -\left(c_t + \boldsymbol{b}_t'\boldsymbol{x}\right)$. Since the normal family is closed under affine transformation, we immediately know the conditional distribution of the loss,

$$L_{t+1}|\mathcal{F}_t \sim \mathcal{N}\left(-c_t - \boldsymbol{b}_t'\boldsymbol{\mu}_{t+1}, \boldsymbol{b}_t'\Sigma_{t+1}\boldsymbol{b}_t\right)$$

hence simple expressions for the VaR and ES.

EWMA

The original model specified zero mean and exponentially-weighted moving average for the covariance with fixed decay parameter θ,

$$\hat{\mu}_{t+1} = \mathbf{0} \qquad \hat{\Sigma}_{t+1} = \theta\boldsymbol{x}_t\boldsymbol{x}_t' + (1-\theta)\hat{\Sigma}_t$$

Historical Simulation Method

A popular non-parametric alternative to the variance-covariance method is historical simulation, which is based on creating an empirical distribution of the loss on the current portfolio by marking it to historical values of the risk factor innovations $\boldsymbol{x}_t, \boldsymbol{x}_{t-1}, \ldots, \boldsymbol{x}_{t-n+1}$ for look-back interval τ_n.

Empirical Loss Distribution

Under this model, the loss distribution function is

$$F_L(l) = \frac{1}{n} \sum_{s=t-n+1}^{t} \chi_{\{l_{[t]}(\boldsymbol{x}_s) \leq l\}}$$

where $\chi_{\{\cdot\}}$ is the indicator function for a logical expression.

VaR and ES are simply the smallest and the average of the largest $k = \lfloor n(1-\alpha) \rfloor + 1$ simulated losses.

Backtesting

Models for measuring risk entail many simplifying assumptions, and both prudence and regulation requires us to monitor our models' performance and our assumptions' validity. Compared to other domains of quantitative finance, the feedback we get from the markets about the quality of risk models is generally subtle.

- For example, if your VaR model says your 99% confidence, ten-day loss is \$10M, and you experienced an actual loss of \$20M, does this mean the model was wrong?

The principal framework we have for monitoring risk model performance is backtesting, which entails comparing time-series of out-of-sample estimates of risk measures with actual realized losses.

Violation-Based Tests

The canonical example is backtesting VaR. The text also includes a discussion of how to go about backtesting ES. A VaR violation is an instance of the event $\left\{L_{t+1} > VaR_t^{(\alpha)}\right\}$. It should be obvious that if $I_{t+1} \triangleq \chi_{\left\{L_{t+1} > VaR_t^{(\alpha)}\right\}}$ is a sequence of random variables on $\Omega = \{0, 1\}$,

$$\mathrm{E}\left[I_{t+1} \mid \mathcal{F}_t\right] = 1 - \alpha$$

and I_{t+1} is a Bernoulli with parameter $1 - \alpha$.
It can be proved that I_t is independent of I_s for $s \neq t$, so this sequence is an i.i.d. process. That means that the sum of such variables in Binomial. It also means that for $\alpha \approx 1$ the spacing between violations is approximately Geometric. This immediately gives us two statistical tests we can apply to the actual VaR violation experience.

Scoring Functions and Elicitability

A very recent (last five years) discussion in measurement has centered on the relevance of the notion of elicitability to risk measures. This comes from the study of performance measurement of forecasts. It is not clear (to me) that risk measures are in fact forecasts; but if they are and this is indeed relevant, we may need to walk back on some of our axioms. In particular, it can be proved that the only coherent elicitable risk measures (except expected loss) are not comonotone additive. But VaR is elicitable.

Scoring Function

The basis of elicitability theory is the existence of a consistent scoring function $S(\cdot,\cdot) \geq 0$ such that

$$\varrho(L) = \arg\min_{y \in \mathbb{R}} \int_{\mathbb{R}} S(y, l)\, dF_L(l)$$

for every possible random variable $L \in \mathcal{M}$.