Value at Risk as a measure of market risk

Susan Thomas
IGIDR, Bombay

16 January 2020
A presentation for the “Knowledge Sharing Session” at RBI, Kolkata

Structure of the talk

- Motivation: financial portfolio risk and capital requirement
- What is Value at Risk (VaR)?
- Calculating VaR
- An example to build intuition
- Volatility forecasting models and where it matters for VaR
- Choosing volatility forecasting models for VaR
- Limits of VaR
- Final thoughts

Financial portfolio risk

Everybody wants to quantify risk

- The trader wants to know much risk she has taken in her trading position.
- The trader's boss, the financial firm, wants to know the risk of each trader, of the firm and how the risk of each trader adds or subtracts from the firm risk.
- The firm's regulators want to know the risk taken by the firm.

How much risk is appropriate?

- If the trader takes no risk, she will obtain no return.
- If the firm asks for very large buffers to protect against the failure of the firm, then it will not make money.
- If the regulator asks firms for very large buffers, to drive down failure probabilities to unreasonably low level, financial activities by regulated firms will subside.
- If we say “SBI should hold enough equity capital to ensure that it will never fail”. In this case, it will need to have 0 borrowing.

The idea of value at risk

- What is a loss so large, that it will be exceeded on 1% of the days?
- This is the “value at risk at a 99% level for a one day horizon”.
- In a year, there *will be* 2-3 days when the loss exceeds the VaR.

Example: VaR at derivatives clearing corporations

- Clearing corporations (CCs) guarantee settlement of exchange traded derivatives trading.
  How?

Example: VaR at derivatives clearing corporations

- Clearing corporations (CCs) guarantee settlement of exchange traded derivatives trading.
  How?
- By collecting a combination of *initial margin* and *daily mark-to-market margin*.
  Latter is easy: Loss in the value of the position at the end of the day.
  Former (initial margin) is difficult $\rightarrow$ big enough to prevent a counterparty from taking flight and small enough to not reduce returns.

Example: VaR at derivatives clearing corporations

- Clearing corporations (CCs) guarantee settlement of exchange traded derivatives trading.
  How?
- By collecting a combination of *initial margin* and *daily mark-to-market margin*.
  Latter is easy: Loss in the value of the position at the end of the day.
  Former (initial margin) is difficult $\rightarrow$ big enough to prevent a counterparty from taking flight and small enough to not reduce returns.
- Initial margin *not* the only fall-back.
  For example, derivatives CCs have • Broker balance sheet, • bank guarantees, additional own capital, • the right to call for additional capital from shareholders, • own capital.
  The ability to guarantee settlement collapses *only after* using up all its own equity capital.
- At the CC, VaR is used to calculate the initial margin such that the one-day MTM loss will be greater than the initial margin on about the 2-3 worst days of the year.

Describing VaR

What is value-at-risk?

The VaR is the loss on a financial portfolio that will be exceeded with a known probability.
In mathematical parlance, if
– the one–day rupee profit on a portfolio is x, and
– returns has a probability distribution function (pdf) $f(x)$, then
– the VaR v at a 95% level is:

$$\int_{-\infty}^{v} f(x)dx = 0.05$$

Inputs to calculating the VaR

1. A known portfolio
   (The VaR of a portfolio is not the sum of the VaR of individual assets).
2. A time horizon of interest.
   Most investors calculate a VaR over a day. Corporations with assets that do not have a daily mark–to–market valuation technology would calculate their VaR over a longer horizon, say, a month.
3. A defined risk level
   A 95% VaR is the loss that will be exceeded on 5 days out of a 100–days.

The VaR is generally reported in rupees.

Interpretation

"The value at risk of this position, on a one-day horizon, at a 99% level, is Rs.1.34 million"

- This means that on 1% of the days, the loss will *exceed* Rs.1.34 million
- This *does not mean* that on the worst 1% of the days, the loss will be roughly Rs.1.34 million.
- In a year of 250 trading days, 1% of the days occur 2-3 days of the time
- So this means: "In the worst 2-3 days of the year, the loss on this position will exceed Rs.1.34 million".

Why is VaR so useful?

- The mathematician says: “I have a full distribution – $f(x;\theta)$. Choose an α and locate the point on the distribution that maps to that level and you know what the expected range of losses can be.”.
- The practical man says: “I cannot understand the full distribution $f(x;\theta)$, but I think I can use the VaR to set capital requirements that can cover losses in a portfolio”.
- VaR is a tool to facilitate conversations between mathematician-statisticians and the finance practitioner.
  The first set know how to calculate it, the second set know how to use it!

Calculating VaR

How to calculate VaR?

- Financial market returns are not i.i.d.
- But let's assume that they are.
- The dgp of financial returns is not stable.
- But let's assume that it is. Let's assume that the future will be like the past.
- These two wrong assumptions gives us a first and simplest way to calculate VaR.

“Historical simulation” (“HS”)

1. Observe the portfolio values for a set of days: calculate changes every day.
2. Sort the data and pick the value that corresponds to 5% worst value.

Example of calculating the 95% VaR using HS

Portfolio: Rs.1,000,000 invested in the equity index for 30 days.
The following are the list of daily changes:

20000619	14415.30	20000703	16045.10	20000718	-11248.10
20000620	5555.85	20000705	9712.49	20000719	-18835.30
20000621	-21427.60	20000706	-6079.85	20000720	-8111.93
20000622	8841.25	20000707	527.29	20000721	-19104.20
20000623	-10843.10	20000710	-5252.30	20000724	-58580.30
20000626	-14057.00	20000711	5878.10	20000725	23363.70
20000627	1823.84	20000712	9698.95	20000726	-22984.30
20000628	10737.70	20000713	-7035.48	20000727	14908.20
20000629	15089.70	20000714	-8475.33	20000728	-3181.32
20000630	-14103.70	20000717	-20138.50	20000731	-712.50

The sorted values and the VaR

- Sort the list of the profits and losses in descending order.

-58580.30	-14103.70	-7035.48	1823.84	10737.70
-22984.30	-14057.00	-6079.85	5555.85	14415.30
-21427.60	-11248.10	-5252.30	5878.10	14908.20
-20138.50	-10843.10	-3181.32	8841.25	15089.70
-19104.20	-8475.33	-712.50	9698.95	16045.10
-18835.30	-8111.93	527.29	9712.49	23363.70

- The 95% VaR is the loss that can be exceeded 5% of the times in this sorted list, is between (a loss of) Rs.22,984 and Rs.21,428.
- The 99% VaR is the loss that can be exceeded 1% of the times in this list, which is a loss of greater than Rs. 58,580.

A parametric model

- A normal distribution for the same returns has a mean close to 0 and standard deviation of 1.46% (daily).
- Then the one-day, 95% VaR for Rs.1,000,000 invested in this asset is:

$$1.96 * 0.0146 * 1000000 = Rs.28,616.00$$

- Problem with this approach:
  Financial market returns are not normally distributed: there are fat tails and volatility clustering.
  Nonlinear products generate non-normal distributions.

A small example to build intuition

Three interesting series

- We look at time series returns for:
  Nifty, INR/USD and one common stock - I T C LTD.
- Daily returns from 1992-01-07 to 2009-12-31, 4700 days of data.

returns

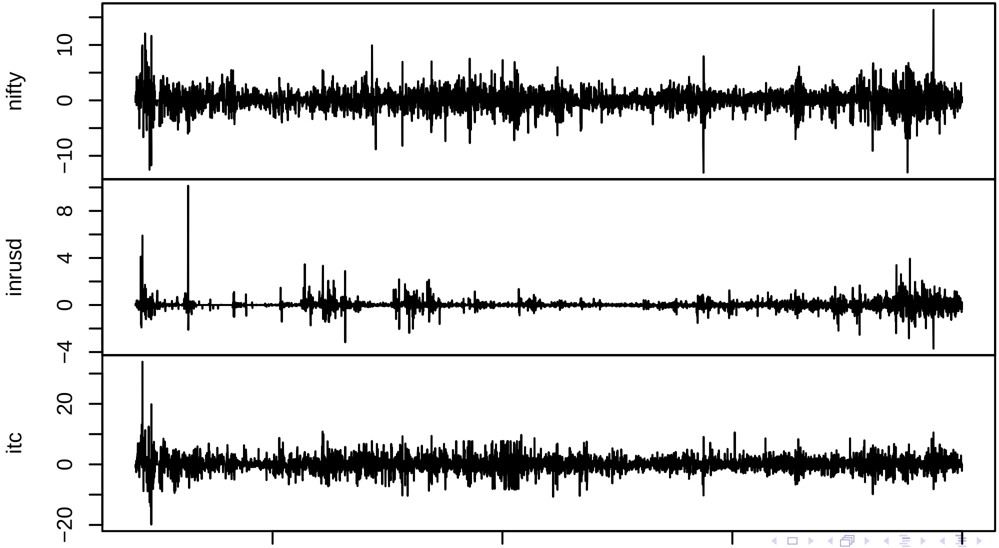

Squared returns

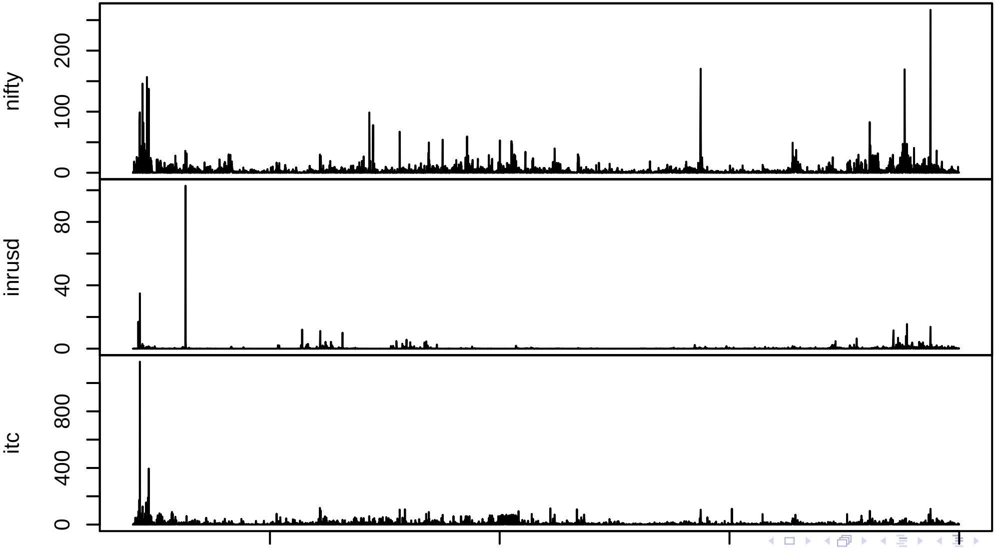

What do we see?

- Volatility clustering!
  Low volatility periods are followed by low volatility periods and vice versa
- There is no “overall average volatility”
- Nifty is less volatile than I.T.C.
- Currency: Most volatile of the three even though the levels are small.
  (But in 2003 and then in 2007 and 2017 things changed).

Nifty kernel density

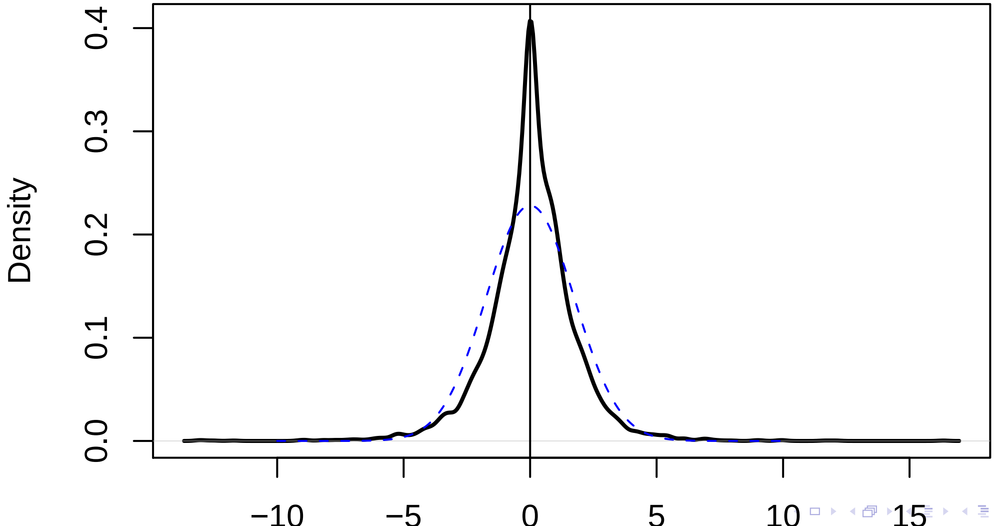

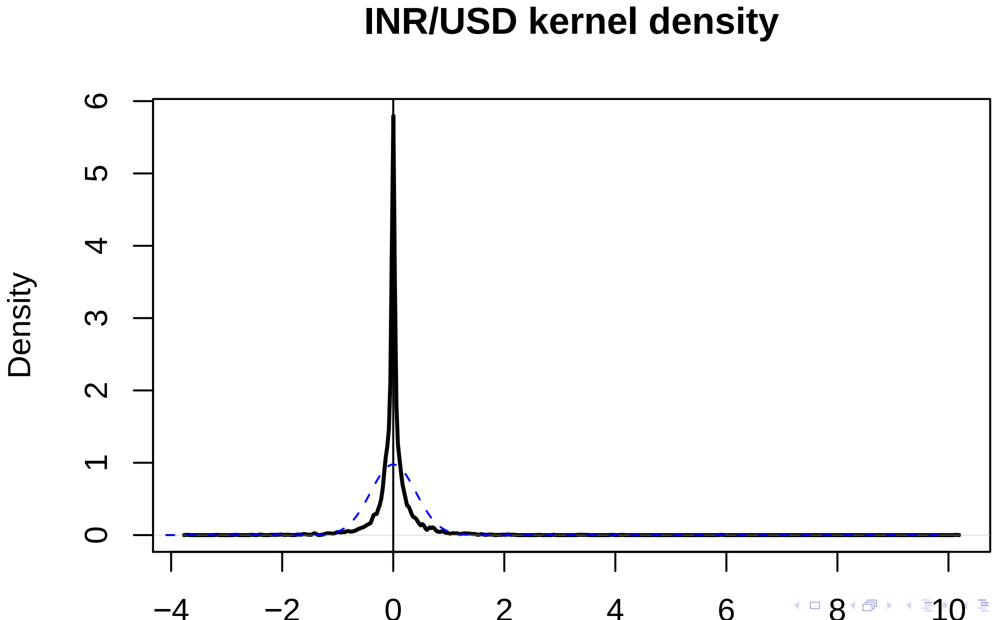
INR/USD kernel density
Density
0
1
2
3
4
5
6
−4
−2
0
2
4
6
8
10

ITC kernel density

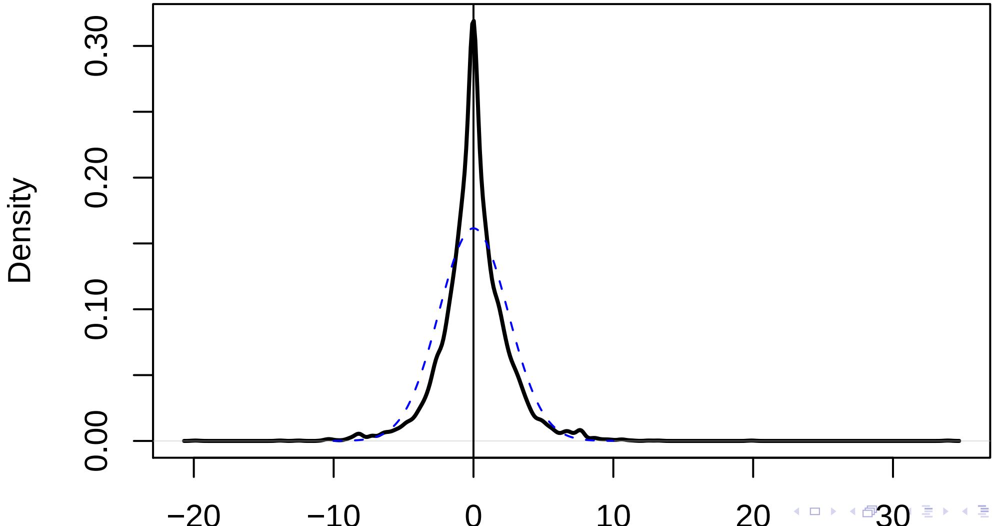

Learnings

- Always have physical intuition about reality – visualise with plots and don't just focus on numbers.
- Move away from assumptions of i.i.d.
- Move away from assumptions about non-normality.
  One way to do this is to build models of heteroscedasticity – constant mean, but changing variance.

Simulations vs. real data

Simulated data: Constant σ

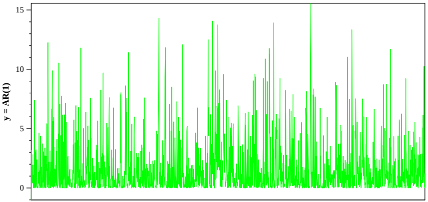

Simulated data: Varying σ

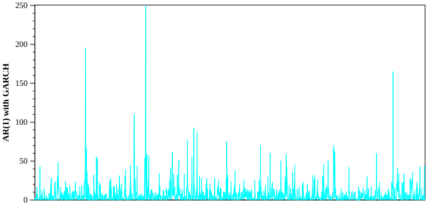

Real data

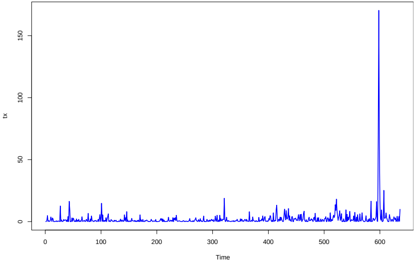

Real data

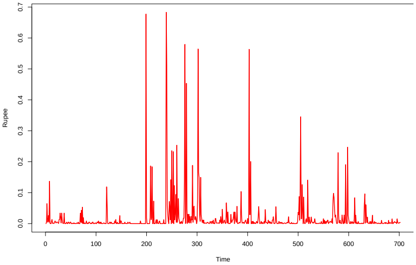

Volatility forecasting models

Varying volatility time series models: ARCH/GARCH

- Wide range of models for dependence in volatility.
- ARCH is the **Auto-Regressive Conditional Heteroskedasticity** (ARCH) model, where ARCH(1) is written as:

$$\begin{aligned} r_t &= \alpha_0 + \alpha_1 r_{t-1} + \epsilon_t; \quad \epsilon_t \sim N(0, \sqrt{h_t}) \\ h_t &= \gamma_0 + \gamma_1 \epsilon_{t-1} \end{aligned}$$

- GARCH is the **Generalised Auto-Regressive Conditional Heteroskedasticity** (GARCH) model where GARCH(1,1) is written as:

$$\begin{aligned} r_t &= \alpha_0 + \alpha_1 r_{t-1} + \epsilon_t; \quad \epsilon_t \sim N(0, \sqrt{h_t}) \\ h_t &= \gamma_0 + \gamma_1 \epsilon_{t-1}^2 + \beta_1 h_{t-1} \end{aligned}$$

Here, h_t is the next period volatility which is a function of the previous volatility (h_{t-1}) and previous period innovation (ϵ_{t-1}).

What works for these models

- They capture the real world data better.
- Even though they are linear models, there are modifications that capture asymmetry (EGARCH) and non-linearity well (regime-switching GARCH).

What works for these models

- They capture the real world data better.
- Even though they are linear models, there are modifications that capture asymmetry (EGARCH) and non-linearity well (regime-switching GARCH).
- They can be extended into multivariate models as well: to model time-series dependence in variance and covariance / correlation jointly.

What works for these models

- They capture the real world data better.
- Even though they are linear models, there are modifications that capture asymmetry (EGARCH) and non-linearity well (regime-switching GARCH).
- They can be extended into multivariate models as well: to model time-series dependence in variance and covariance / correlation jointly.
- The distribution of returns after adjusting for variance dependence *tend* to be more normal.
- They forecast volatility better than historical variance
  $\rightarrow$ better VaR estimates
  $\rightarrow$ better estimates of capital requirements for financial firms.

What does not work for these models

- Too many choices!
- The translation into multi-variate forms is not easy.
  Multivariate GARCH models often need simplifications to make estimations tractable (DCC-GARCH models).
- These models are linear and deterministic: volatility may have a stochastic nature.

Recent developments in volatility forecasts for VaR

- Stochastic volatility models (SV)
- Models to fit the tail behaviour of financial asset returns.
  This includes
  1. Extreme Value Theory and Copulas – as distributions for just the tail – and
  2. Expected Loss models to capture the distribution beyond the VaR point estimate.
- Use of high frequency data to capture intra-day dependence, which is then aggregated to create new models for volatility forecast.
- Use of asynchronous data – move from constant time periodicity to constant volume periodicity to model dependence.
- Use of market data to directly observe volatility – *Implied Volatility* estimates from options markets.

Problem: How to choose which fits *best* for VaR?

Evaluating model performance using VaR

Testing which volatility model fits the world best

- Simple idea: we use the volatility models to forecast VaR so that we can get the correct capital requirement.
  When we choose $\alpha = 95\%$, then daily VaR should be exceeded 5% of the time, i.e. about 18 days a year.
  Question: Does it do so?
- The volatility forecasting model where the VaR is exceeded at the “correct” α number of times is the “best” forecasting model.

Volatility model performance evaluation using VaR

- Suppose there are two forecasting models which give two forecasts of VaR (V_1, V_2).
- Compare (V_1, V_2) against actual return, r_{t+1}.
  If $V_1 < r_{t+1}$ and $V_2 > r_{t+1}$, then
  V_1 is *better* than V_2.

Volatility model performance evaluation using VaR

- Suppose there are two forecasting models which give two forecasts of VaR (V_1, V_2).
- Compare (V_1, V_2) against actual return, r_{t+1}.
  If $V_1 < r_{t+1}$ and $V_2 > r_{t+1}$, then
  V_1 is *better* than V_2.
- Statistical test: Repeat this many times, and see how many times $V_1 < r$ and $V_2 < r$.
  Which comes closer to p is the better model.
- Real world complications: but there should be no clustering!
- Real world complications: in the above test, the magnitude of difference between V_t and r_t doesn't matter. But we care about how much this difference is as banking supervisor.

Tests of forecast performance using VaR

- Performance evaluation using VaR #1: Statistical *unconditional coverage* and *conditional coverage* tests by Christoffersen (1998).
  Transform the data on r_t, v_t into I_t, where:
  1. $I_t = 1$ if $v_t > r_t$, and
  2. $I_t = 0$ if not.

  The forecasts are efficient if they show both correct unconditional coverage and no independence.

Tests of forecast performance using VaR

- Performance evaluation using VaR #1: Statistical *unconditional coverage* and *conditional coverage* tests by Christoffersen (1998).
  Transform the data on r_t, v_t into I_t, where:
  1. $I_t = 1$ if $v_t > r_t$, and
  2. $I_t = 0$ if not.

  The forecasts are efficient if they show both correct unconditional coverage and no independence.
- Performance evaluation #2: Loss function approach by Lopez (1998).
  Calculate l_t as:
  1. Example 1: $l_t = (r_t - v_t)^2$ if $v_t > r_{t+1}$; 0 if not.
  2. Example 2: $l_t = (r_t - v_t)^2$ if $v_t > r_{t+1}$; $-\alpha v_t$ if not.

  In Example 2, the larger the error, the greater the penalty to the model.
- Sarma et al (2003) used both these tests to find the "best" volatility forecasting model for the Indian equity index to set the initial margin for derivatives trading.

A technical digression

Desirable properties for risk measures

Purpose: to summarise the entire distribution of value returns X by a single measure.

Desirable properties for risk measures

Purpose: to summarise the entire distribution of value returns X by a single measure. Features:

- 1. Monotonicity: if $X_1 \leq X_2$, then $RM(X_1) \leq RM(X_2)$

Desirable properties for risk measures

Purpose: to summarise the entire distribution of value returns X by a single measure. Features:

- 1. Monotonicity: if $X_1 \leq X_2$, then $RM(X_1) \leq RM(X_2)$
- 2. Translation invariance: $RM(X + C) = RM(X) + C$ where C is cash. Higher levels of cash means lower the risk of the portfolio.

Desirable properties for risk measures

Purpose: to summarise the entire distribution of value returns X by a single measure. Features:

- 1. Monotonicity: if $X_1 \leq X_2$, then $RM(X_1) \leq RM(X_2)$
- 2. Translation invariance: $RM(X + C) = RM(X) + C$ where C is cash. Higher levels of cash means lower the risk of the portfolio.
- 3. Homogeneity: $V(mX) = mV(X)$
  A portfolio with the same constituents that is larger by a multiplier m will have a higher risk value, scaled up by multiplier m.

Desirable properties for risk measures

Purpose: to summarise the entire distribution of value returns X by a single measure. Features:

- 1. Monotonicity: if $X_1 \leq X_2$, then $RM(X_1) \leq RM(X_2)$
- 2. Translation invariance: $RM(X + C) = RM(X) + C$ where C is cash. Higher levels of cash means lower the risk of the portfolio.
- 3. Homogeneity: $V(mX) = mV(X)$
  A portfolio with the same constituents that is larger by a multiplier m will have a higher risk value, scaled up by multiplier m.
- 4. Subadditivity: $V(X_1 + X_2) \leq V(X_1) + V(X_2)$
  Combining portfolios cannot increase risk.

Is VaR a “good measure”?

- VaR fits the criteria of monotonicity and translation invariance (#1 and #2).

Is VaR a “good measure”?

- VaR fits the criteria of monotonicity and translation invariance (#1 and #2).
- On homogeneity (#3):
  VaR is calculated on returns distribution of the portfolio.
  But large portfolios have a larger price impact. If this is not taken into account, VaR does not satisfy the homogeneity property.

Is VaR a “good measure”?

- VaR fits the criteria of monotonicity and translation invariance (#1 and #2).
- On homogeneity (#3):
  VaR is calculated on returns distribution of the portfolio.
  But large portfolios have a larger price impact. If this is not taken into account, VaR does not satisfy the homogeneity property.
- Subadditivity of credit portfolios may report VaR as having a higher VaR than an investment in any single credit instrument.

Is VaR a “good measure”?

- VaR fits the criteria of monotonicity and translation invariance (#1 and #2).
- On homogeneity (#3):
  VaR is calculated on returns distribution of the portfolio.
  But large portfolios have a larger price impact. If this is not taken into account, VaR does not satisfy the homogeneity property.
- Subadditivity of credit portfolios may report VaR as having a higher VaR than an investment in any single credit instrument.

VaR can be dangerous is the simplicity: a single measure that captures *the edge* of how bad things can get. But does not tell us *everything* about how bad things can get.

Final thoughts

VaR is an approximation with several assumptions

1. Distribution of the returns
2. Dependence on observed data and patterns – backward looking bias towards most recent observations.
3. Pro-cyclicality
4. Poor at dealing with structural change
5. Ease of use may not be a good thing – standardising risk perceptions may not be the best thing from a "systemic risk" point of view (regulatory bias).
6. Does not explicitly incorporate different factors of risk: market liquidity, credit.
7. Without a full set of markets to lay off different types of risks, a VaR measurement is ineffective.
8. An illusion of safety: pushing α from 99.9999% to 99.999999% gets very little additional safety and can be expensive to the financial sector.

All these are areas where innovative research can be useful to provide feedback and corrections on how to measure and use VaR better.

Learnings for financial sector risk management

- A fairly good VaR estimate is possible.
- But marking to market is the first required step.
  Without a “current” value of the portfolio, volatility prediction is rendered ineffective.

Learnings for financial sector risk management

- A fairly good VaR estimate is possible.
- But marking to market is the first required step.
  Without a "current" value of the portfolio, volatility prediction is rendered ineffective.
- In most liquid markets of the world (e.g. the US equity market, the global currency market), there are several problems with using VaR.
  In the less liquid markets (e.g. financial markets in India, e.g. credit risk) these problems can be exacerbated.
  If we cannot get a current value, we can neither measure risk nor manage it well.

Learnings for financial sector risk management

- A fairly good VaR estimate is possible.
- But marking to market is the first required step.
  Without a "current" value of the portfolio, volatility prediction is rendered ineffective.
- In most liquid markets of the world (e.g. the US equity market, the global currency market), there are several problems with using VaR.
  In the less liquid markets (e.g. financial markets in India, e.g. credit risk) these problems can be exacerbated.
  If we cannot get a current value, we can neither measure risk nor manage it well.
- Illiquid securities should have higher capital requirements since uncertainty about their value is high
- An optimal path to risk management will be to focus policy on making markets more liquid.
  With liquid markets, we can get better ("current") values.
  With current values, the Bank can manage risk better.

Thank you.

susant@igidr.ac.in

www.ifrogs.org