Lecture notes on risk management, public policy, and the financial system

Value-at-Risk

Allan M. Malz

Columbia University

Why Value-at-Risk?

- Value-at-Risk (VaR) of a portfolio: single number summarizing risk of large and complex portfolios
  - "How much can we lose?" You can't refuse to answer!
  - Encompasses different asset types
- Reasonably accurate for many types of portfolios
  - Unusual but recurrent losses, not extremes
- **VaR limit system:** position size limits based on VaR
  - Form of **risk budgeting**
  - Widely-used to control risk while giving some discretion to individual trading desks
- Can be computed using broad range of return models, estimation method, data sources
- Most importantly, we can learn a lot from criticizing it!

Definition: VaR is a quantile

- VaR of a portfolio: a **quantile** of the portfolio **loss distribution**
  - Loss distribution: $-1\times$ profit-and-loss (P&L) distribution
  - Loss and VaR defined as *positive* numbers, in dollar or return units
- p-quantile of a random variable (r.v.) X:

$$X^\circ \text{ s.t. } \mathbf{P}\left[X \leq X^\circ\right] = p$$

  - The value X° with cumulative probability p
  - Threshold X° below which realizations of X fall with frequency p
- To define VaR, let X represent the r.v. loss distribution, and α the **confidence level** of the VaR estimate
- VaR at confidence level α is α-quantile of loss distribution
  - Probability of losing no more than the VaR is α, e.g., 99 percent
  - Probability of suffering loss worse than VaR is $1-\alpha$, e.g., 1 percent
- Daily VaR at 95 (99) percent confidence level should occur roughly one in 20 trading days, or once per month (twice a year)

Example of a quantile and VaR

- With $\Phi(X)$ the standard normal c.d.f., $z_p = \Phi^{-1}(p)$ is the standard normal **inverse cumulative distribution** or **quantile function**
- **Example:** next-day level of S&P 500 as of 28Aug2013
- Changes in S&P lognormally distributed$\Rightarrow$ S&P log return normally distributed
    - Quantiles of S&P log return $=$ standard normal quantiles $\times$ estimated return volatility

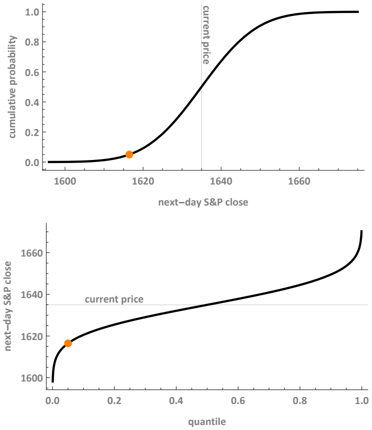

Common misconceptions about VaR

- VaR wedded to normally-distributed return model
- Or to a particular way of using market data
- And some outright distortions, e.g. “VaR is the most I can lose”
- Complex pros and cons for each judgement call in VaR modeling
  - But need for judgement calls in risk modeling not unique to VaR

The VaR scenario

- The **VaR scenario** is the quantile of P&L corresponding to the chosen confidence level
- Can be stated as P&L (in currency units) or as adverse return (decimal or percent)
- Models generally based on distributional hypothesis about log returns
  - VaR scenario stated as log return r° corresponds to P&L $xS_t(e^{r^{\circ}} - 1)$
  - VaR scenario stated as arithmetic return $r^{\text{arith},\circ} = e^{r^{\circ}} - 1$ corresponds to same P&L $xS_t r^{\text{arith},\circ}$

VaR example

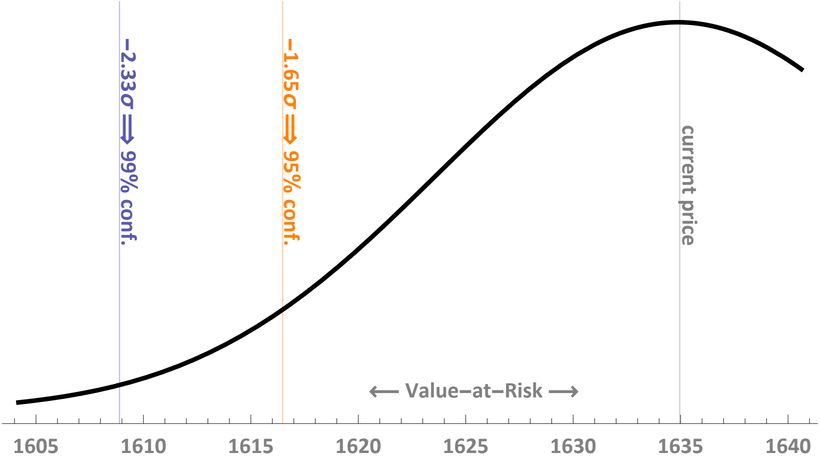

1-day VaR of a long S&P 500 index position on 28Aug2013. Volatility computed via EWMA with a decay factor of 0.94. Grid lines at VaR scenarios for confidence levels of 95 and 99 percent. VaR is the difference between the index value in the VaR scenario and 1634.96, its 28Aug2013 closing value, times the number of index units held.

How to compute VaR

Three basic approaches to identifying the VaR scenario:

Parametric is a simple approach relying on a formula based on a hypothesized return distribution plus a volatility estimate

Monte Carlo simulation uses random draws from a hypothesized return distribution

- Returns an output of pricing models rather than direct observations
- Used for portfolios that include derivatives, complex securities, e.g. MBS

Historical simulation is based on historical returns over some past observation period, no distributional hypothesis

Judgments and data decisions in VaR estimation

Distributional hypothesis: What asset return model: normal, lognormal, *t*-distribution,...? How estimate parameters: GARCH, EWMA,...?

Risk factor mapping: VaR generally applied with small number of **risk factors** relative to number of positions in portfolio

- Most assets' market risks more accurately modeled as functions of factor rather than own-price risks
- E.g. equity risks function of index, Fama-French factors, bonds a function of key points on interest-rate curve
- Reduces computational complexity
- To which risk factors is portfolio exposed? Many pitfalls, for example:
  - Mapping AAA subprime mortgage bonds to AAA corporate bonds
  - Omitting key risk factor, such as option implied volatility for option portfolio

Use of historical data: How much history? Include or exclude extreme and possibly "unique" events?

User settings in VaR modeling

- User makes decisions based on business application about:
  Time horizon τ over which "worst-case" P&L realized
  Confidence level α that losses will be no worse than VaR
- VaR is generally *higher* at longer time horizons and higher confidence levels
- VaR is generally *less accurate* at longer time horizons and higher confidence levels
- Problematic: setting (→)**economic capital** based VaR
  - Capital should be set high enough to cover rare, but large and costly, losses
  - But VaR more accurate at predicting recurrent losses at "cost-of-doing-business" level
  - VaR can be interpreted as maximum loss if extreme event *does not* occur
- VaR generally treated as one-tailed test↔one tail corresponds to losses
  - Exceptions include option portfolios
  - And not necessarily the left tail (→short positions)

Typical model assumptions for VaR

- Logarithmic asset price changes $r_{t,t+\tau} \equiv \ln(S_{t+\tau}) - \ln(S_t)$ normally distributed
  - Including zero-mean assumption
    $$r_{t,t+\tau} \sim \mathcal{N}(0, \sigma_t^2 \tau)$$
- Volatility estimate σ_t based on information up to time t but constant over any future horizon τ
  - $\Rightarrow$Use square-root-of-time rule to apply volatility estimate to any horizon
- For confidence level α, 1-day horizon, and for long position in the risk factor/asset, take $1-\alpha$ quantile of $r_{t,t+\tau}$
- Parametric: $z_{1-\alpha}$ quantile of standard normal
  - Typically a negative number
  - Log return in the VaR scenario estimated as $z_{1-\alpha}\sigma_t\sqrt{\tau}$
- Monte Carlo: $1-\alpha$ quantile of simulations of r_t

VaR computation example: data and assumptions

- Calculation date 28Aug2013
- Risk factor S&P 500 index, closed at $S_t = 1634.96$
- **Exposure**: long position with initial value $xS_t =$\$1 000 000, with $x \equiv$number of units of asset
  - $\Rightarrow x = \frac{1\,000\,000}{1634.96} = 611.636$ index units
- One-tailed test at confidence level 99 percent
  - Corresponding standard normal quantile $z_{0.01} = -2.32635$
- Note that computation doesn't require x and S_t individually, just initial position value xS_t and return quantile
- σ_t estimated at close on 28Aug2013 via EWMA, with $\lambda = 0.94$
  - Pertains to any future horizon using square-root-of-time rule
  - Volatility estimate on 28Aug2013 $\sigma_t = 0.0069105$ or 69 bps/day
    - Annualized vol about 11.06 percent, relatively low for S&P
  - Used in computing VaR parametrically and via Monte Carlo, not via historical simulation
- One-day horizon: $\tau = 1$, with time measured in days, volatility at daily rate

Parametric normal VaR: theory

VaR scenario in log return terms: at 99-percent confidence level, use 0.01-quantile of $r_{t,t+\tau}$ $z_{1-\alpha}\sigma_t\sqrt{\tau}$, the $1-\alpha$ quantile of $r_{t,t+\tau}$

- Our model tells us the log return is normal

Change in risk factor in VaR scenario: convert log return into arithmetic return needed to compute P&L

- $1-\alpha$ quantile of $S_{t+\tau} - S_t$ is:

$$S_t e^{z_{1-\alpha}\sigma_t\sqrt{\tau}} - S_t = S_t\left(e^{z_{1-\alpha}\sigma_t\sqrt{\tau}} - 1\right)$$

VaR scenario in P&L terms: $1-\alpha$ quantile of change in position value is $xS_t\left(e^{z_{1-\alpha}\sigma_t\sqrt{\tau}} - 1\right)$

- Multiplies quantile of change in risk factor by position size x

VaR at confidence level α is P&L quantile expressed as positive number:

$$\mathrm{VaR}_t(\alpha,\tau) = -xS_t\left(e^{z_{1-\alpha}\sigma_t\sqrt{\tau}} - 1\right) = xS_t\left(1 - e^{z_{1-\alpha}\sigma_t\sqrt{\tau}}\right)$$

Parametric normal VaR: example

VaR scenario in log return terms: a bit over -1$\frac{1}{2}$ percent

$$z_{0.01}\sigma_t\sqrt{1} = -2.32635 \times 0.0069105 = -0.0160762$$

Change in risk factor in VaR scenario: 0.01-quantile of 1-day change $S_{t+1} - S_t$:

$$S_t\left(e^{z_{0.01}\sigma_t} - 1\right) = 1634.96(e^{-0.0160762} - 1) = 1634.96 \times (-0.0159477)$$
$$= 1608.89 - 1634.96 = -26.07$$

VaR scenario in P&L terms:

$$xS_t\left(e^{z_{0.01}\sigma_t} - 1\right) = 1\,000\,000 \times (-0.0159477)$$

VaR at a 99-percent confidence level is \$15 947.70:

$$xS_t\left(1 - e^{z_{0.01}\sigma_t}\right) = 1\,000\,000 \times 0.0159477 = 15\,947.66$$

- 1-week (5 business days) VaR is \$35 309.00:

$$1\,000\,000\left(1 - e^{0.0160762\sqrt{5}}\right) = 1\,000\,000 \times 0.035309$$

A convenient approximation of VaR

- Apply $e^a - 1 \approx a$ to pertinent quantile:

$$-S_t z_{1-\alpha}\sigma_t\sqrt{\tau} \approx S_t\left(1 - e^{z_{1-\alpha}\sigma_t\sqrt{\tau}}\right)$$

  - Use slightly smaller (larger loss) log return in place of arithmetic return
- Tantamount to assuming arithmetic—not log—returns normally distributed
  - And treats σ_t as estimate of volatility of arithmetic returns
- Approximation widely used, e.g. in (→)**delta-normal** approach to VaR computation
- VaR can also be expressed in return terms, i.e. ≈1.6 percent rather than ≈\$16 000
- Approximation in our **example** using vol at daily rate:
  - 1-day VaR: $-xS_t z_{0.01}\sigma_t = \$16\,076.20$
  - 5-day VaR: $-xS_t z_{0.01}\sigma_t\sqrt{\tau} = \$35\,947.50$

Monte Carlo computation of VaR

- Steps in the algorithm:
  1. Generate set of, say, 10 000 independent draws $\epsilon_i, i = 1, \ldots, 10\,000$ from standard normal
  2. Each draw provides a random realization of log return $r_{t+1}, r_i = \sigma_t \epsilon_i$, next-period price S_i, position value xS_i, and P&L $x(S_i - S_t)$
  3. Sort the realizations in ascending order (largest loss first)
     - $\rightarrow$**order statistics** $\tilde{r}^{(i)}, \tilde{S}^{(i)} - S_t = S_t \left(e^{\tilde{r}^{(i)}} - 1 \right)$ or $\tilde{S}^{(i)} - S_t \approx \tilde{r}^{(i)} S_t$
  4. The 100th order statistic of P&L $\times(-1)$ corresponds to the VaR at a 99-percent confidence level
- Monte Carlo requires estimate of volatility and other model parameters
  - In our example, we've posited lognormal/zero-drift model
  - But as with parametric, no particular model required
- Result \$15 912.46 or thereabouts
  - User may average or interpolate scenarios near the VaR (in our example, near the 100th) to reduce simulation noise

Monte Carlo computation of VaR: example

i	$\tilde{r}^{(i)}$	$\tilde{S}^{(i)}$	P&L
1	-0.02552	1593.76	-25,198.23
2	-0.02517	1594.33	-24,851.55
3	-0.02477	1594.96	-24,465.89
⋮	⋮	⋮	⋮
99	-0.01606	1608.92	-15,928.02
100	-0.01604	1608.94	-15,912.46
101	-0.01602	1608.97	-15,893.97
⋮	⋮	⋮	⋮
4999	0.00005	1635.04	50.42
5000	0.00005	1635.05	53.70
5001	0.00006	1635.05	57.12
⋮	⋮	⋮	⋮
9998	0.02344	1673.73	23,713.88
9999	0.02351	1673.86	23,793.06
10000	0.02407	1674.80	24,365.89

Entries in the second column are $\tilde{r}^{(i)} = \sigma_t \tilde{\epsilon}^{(i)}$, where the $\tilde{\epsilon}^{(i)}$ are the ordered draws from $\mathcal{N}(0, \sigma_t^2)$, with $\sigma_t = 0.0060762$. Entries in the second column are $\tilde{S}^{(i)} = S_t e^{\tilde{r}^{(i)}}$. The P&L realizations are $10^6 \left(e^{\tilde{r}^{(i)}} - 1\right)$.

Monte Carlo computation of VaR

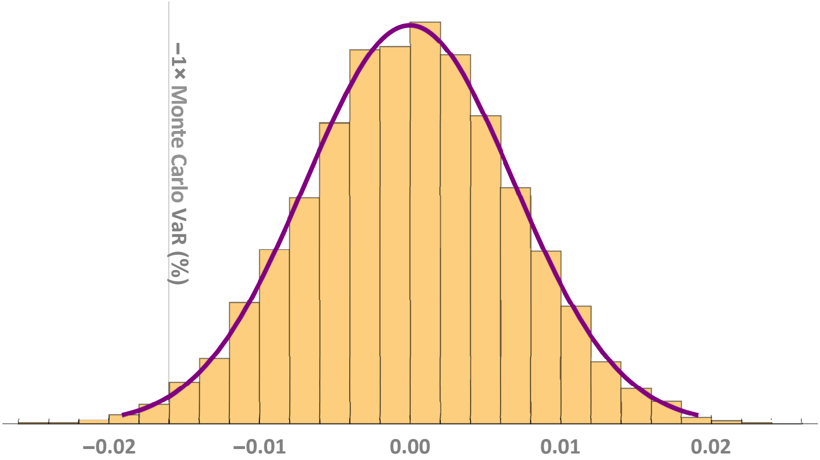

Histogram of Monte Carlo return simulations. Purple plot: density of $\mathcal{N}(0, \sigma_t^2)$, with $\sigma_t = 0.0068826$.

Steps in the algorithm

1. Select a historical "look-back" period, say, 2 years, and compute $t = 1, \ldots, m$ daily log or arithmetic returns
   - Use historical risk factor returns but current portfolio position sizes or weights
2. From here, procedure identical to Monte Carlo: sort m historical realizations in ascending order
   - Order statistics denoted $\tilde{r}^{(i)}$ or $\tilde{r}^{\text{arith},(i)}, i = 1, \ldots, m$
3. Use the $\tilde{r}^{(i)}$ or $\tilde{r}^{\text{arith},(i)}$ to get m ordered simulations of P&L:

$$\tilde{r}^{\text{arith},(i)} x S_t = \left(e^{\tilde{r}^{(i)}} - 1\right) x S_t, \qquad i = 1, \ldots, m$$

4. VaR of long position at confidence level α is $\times(-1)$ the $(1-\alpha)$-quantile of order statistics of P&L

Quantiles of empirical distributions

- Different definitions of quantile can lead to different results
- General definition of p-quantile of the distribution of an r.v. X

$$X^\circ = \inf\{X | \mathbf{P}[X \leq X^\circ] \geq p\}$$

  - p-quantile X° is *smallest* value of X s.t. the cumulative probability of X° is *at least p*
- Applied to VaR at confidence level α: smallest loss X° such that the probability of a loss no larger than X° is at least α
- Definition applies to both continuous (e.g. normal) and discrete distributions (e.g. simulations)
  - Discrete distributions if $(1-\alpha)m$ an integer
- But leads to unambiguous result only for
  - Continuous distributions that are not flat at $1-\alpha$
  - Discrete distributions if $(1-\alpha)m$ an integer

Identifying quantiles of empirical distributions

- General definition consistent with many alternative methods for empirical distributions
- Which order statistic $\tilde{r}^{(i)}$ represents $(1-\alpha)$-quantile?
  - Most commonly-used is **ceiling** $\lceil(1-\alpha)m\rceil$ of $(1-\alpha)m$: smallest integer $\geq (1-\alpha)m$
  - Or **floor** $\lfloor(1-\alpha)m\rfloor$ of $(1-\alpha)m$: largest integer $\leq (1-\alpha)m$
  - Or interpolate between $\lfloor(1-\alpha)m\rfloor$-th and $\lceil(1-\alpha)m\rceil$-th order statistics
- These methods lead to same result if $(1-\alpha)m$ an integer

Choosing the VaR scenario via historical simulation

- Typically fewer simulations when using historical rather than computer-generated simulations
- $\rightarrow$Potentially material sensitivity of historical simulation result to choice of quantile definition
- Definition of cumulative probability is asymmetrical: event for which probability is defined is $X \leq X^{\circ}$
  - Random variables are **right-continuous**
  - Quantile function therefore **left-continuous**

Example of historical simulation VaR

- \$1 000 000 long position in S&P 500 index
- Using 2 years of price data, 28Aug2011 to 28Aug2013
  - $m = 503$ return observations
- VaR at 99-percent confidence level is \$26 705 using most-common quantile definition, much higher than parametric or Monte Carlo
  - $(1-\alpha)m = 0.01 \cdot 503 = 5.03$
  - $\lceil (1-\alpha)m \rceil = 6$
  - -0.02671 is 6th order statistic of arithmetic returns $\tilde{r}^{\text{arith},(i)}, i = 1, \ldots, m$ in the historical sample
  - 0.02671 is the smallest loss s.t., within the set of return observations, the frequency of a loss no greater than that is at least 99 percent
- VaR can be reported as lying elsewhere on interval $\left[\tilde{r}^{(5)}, \tilde{r}^{(6)}\right]$ bet. 5th and 6th P&L order statistics using alternative quantile definitions

Order statistics for historical simulation VaR

i	t	S_t	date t	S_{t-1}	$\tilde{r}^{(i)}$	$\tilde{r}^{\text{arith},(i)}$	P&L
1	52	1229.10	09Nov2011	1275.92	-0.03739	-0.03670	-36 695.09
2	18	1129.56	22Sep2011	1166.76	-0.03240	-0.03188	-31 883.16
3	17	1166.76	21Sep2011	1202.09	-0.02983	-0.02939	-29 390.48
4	25	1099.23	03Oct2011	1131.42	-0.02886	-0.02845	-28 450.97
5	46	1218.28	01Nov2011	1253.30	-0.02834	-0.02794	-27 942.23
6	9	1154.23	09Sep2011	1185.90	-0.02707	-0.02671	-26 705.46
7	5	1173.97	02Sep2011	1204.42	-0.02561	-0.02528	-25 281.88
8	455	1588.19	20Jun2013	1628.93	-0.02533	-0.02501	-25 010.28
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
499	64	1192.55	28Nov2011	1158.67	0.02882	0.02924	29 240.42
500	80	1241.31	20Dec2011	1205.35	0.02940	0.02983	29 833.66
501	30	1194.89	10Oct2011	1155.46	0.03356	0.03412	34 124.94
502	43	1284.59	27Oct2011	1242.00	0.03372	0.03429	34 291.47
503	66	1246.96	30Nov2011	1195.19	0.04240	0.04332	43 315.29

The entries in the last 3 columns are the order statistics of the logarithmic and arithmetic historical return and P&L realizations: $\tilde{r}^{(i)}$, $e^{\tilde{r}^{(i)}} - 1$, and $xS_m\left(e^{\tilde{r}^{(i)}} - 1\right)$, $i = 1, \ldots, m$ and $m = 503$. The VaR scenario is highlighted.

Computation of VaR by historical simulation

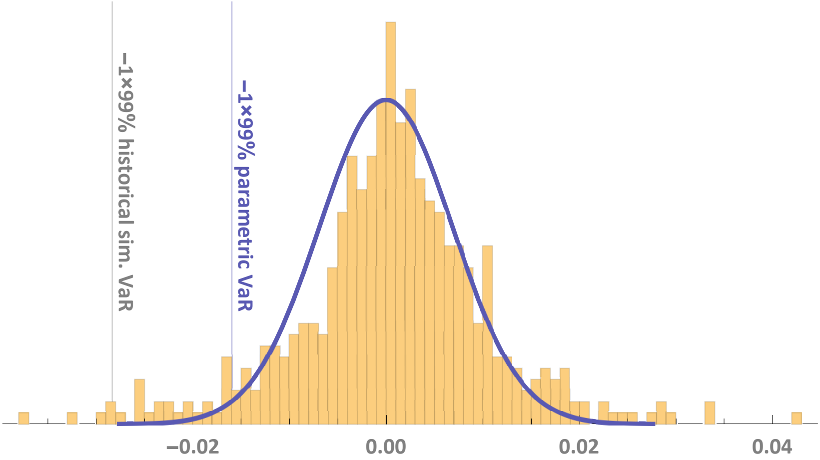

Histogram of historical returns. Purple plot: density of $\mathcal{N}(0, \sigma_t^2)$, with $\sigma_t = 0.00691049$.

Historical simulation VaR scenario

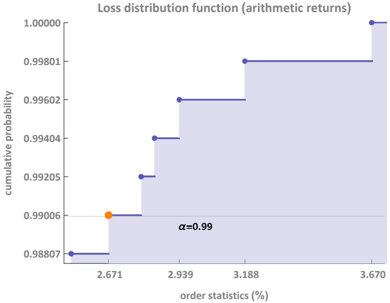

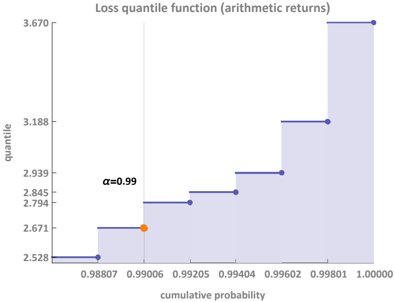

Purple points identify $(-1)\times$arithmetic return observations in the left tail. These are the magnitudes of the return observations leading to the largest observed losses in the historical sample. Orange point denotes quantile using $\lceil (1-\alpha)m \rceil$-th order statistic. With $\alpha = 0.99$ and $m = 503$, $\lceil (1-\alpha)m \rceil = 6$, and the VaR scenario at a 99-percent confidence level in arithmetic return terms is -2.671 percent.

VaR for short positions

- Definition of VaR unchanged: low quantile of P&L
- But $x < 0 \Rightarrow$VaR return scenario positive, not negative
  - Scenario in upper tail of return distribution
- Major drawback of VaR for short positions: doesn't capture unlimited downside
  - P&L of short$\rightarrow -\infty$ as $S_{t+\tau} \rightarrow \infty$

Normal parametric VaR for short positions

- For confidence level α, use α- rather than $(1-\alpha)$-quantile
- VaR log return scenario estimated at t is $z_\alpha \sigma_t \sqrt{\tau} > 0$
- VaR at confidence level α: P&L quantile $xS_t \left(e^{z_\alpha \sigma_t \sqrt{\tau}} - 1\right)$ expressed as a positive number:

$$\mathrm{VaR}_t(\alpha, \tau) = (-1) \times xS_t \left(e^{z_\alpha \sigma_t \sqrt{\tau}} - 1\right)$$

- Normal is a symmetric distribution$\Rightarrow z_\alpha = -z_{1-\alpha}$
- VaR slightly larger than for long position, since $e^r - 1 > 1 - e^{-r}$
- Approximation $xS_t z_\alpha \sigma_t$ gives same value as for long
- Continuing the **example:** $\mathrm{VaR}_t(0.99, 1) = \$16\,206.10$

Historical simulation VaR for short positions

- Basic simulation approach unchanged
- Several equivalent ways to identify VaR scenario in return terms
  - Use a high (i.e. α) quantile of historical return series
  - Use low $(1-\alpha)$ order statistic of $(-1)\times$ historical returns
  - Sort the returns in reverse order, and then use the rank corresponding to a low quantile

Pitfalls in using less-common quantile definitions

- Arise from asymmetric definition of distribution function
- **Example**: historical simulation VaR of long and short S&P 500 position, $\alpha = 0.99$ and $m = 503$
  - Linear interpolation using order statistic of $(-1)\times$ original series between $\tilde{r}^{(5)}$ and $\tilde{r}^{(6)}$, the 5th and 6th smallest
  - Linear interpolation using original order statistics between $\tilde{r}^{(m-5)}$ and $\tilde{r}^{(m-6)}$, the 6th and 7th largest

Historical simulation VaR for short position

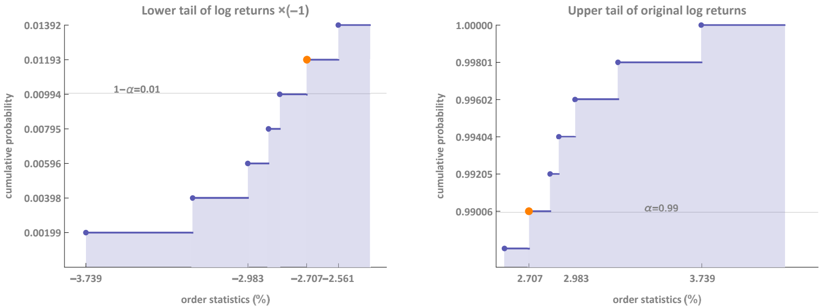

Purple points denote log return observations in the tails. Orange points denote quantiles using $\lceil (1-\alpha)m \rceil$-th order statistic.

Capturing the tails

- Monte Carlo almost identical to parametric for simple/linear portfolios
  - Simulations merely reflect the simulated distribution
  - Monte Carlo becomes useful in more complex portfolios (options, other non-linear assets)
    - Simulated risk factor returns become inputs into pricing models
- Historical simulation may differ greatly from Monte Carlo or parametric
  - Historical simulations may have thicker or thinner tails than Monte Carlo or parametric
  - Depends on length of historical look-back period
  - How far back should we look?
    - Depends on purpose of estimate: recurrent losses or extreme events?
    - How to treat the period mid-2007 to date?

Comparison of VaR computation approaches

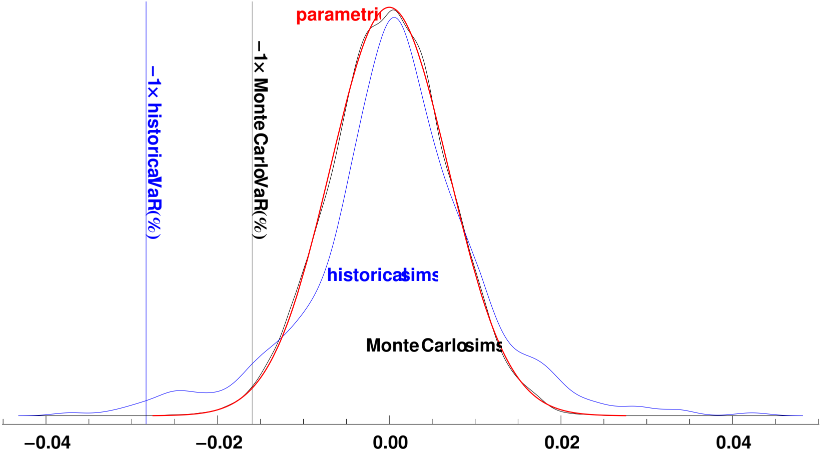

Computed for S&P 500 index on 28Aug2013. Kernel density estimates for Monte Carlo simulations (black) and of historical returns (red). Color-coded vertical grid lines placed at 0.01 quantiles of each distribution.

Incorporating conditionality

- Parametric and Monte Carlo VaR much more responsive to recent returns than historical simulation
- Sluggish responsiveness of historical simulation mitigated by shorter look-back period
- Historical simulations may have thicker or thinner tails than Monte Carlo or parametric
  - Shorter observation intervals may miss tail events
  - Longer observation intervals may produce results deviating from current return distribution (e.g. **volatility regime**)

VaR responsiveness to shocks

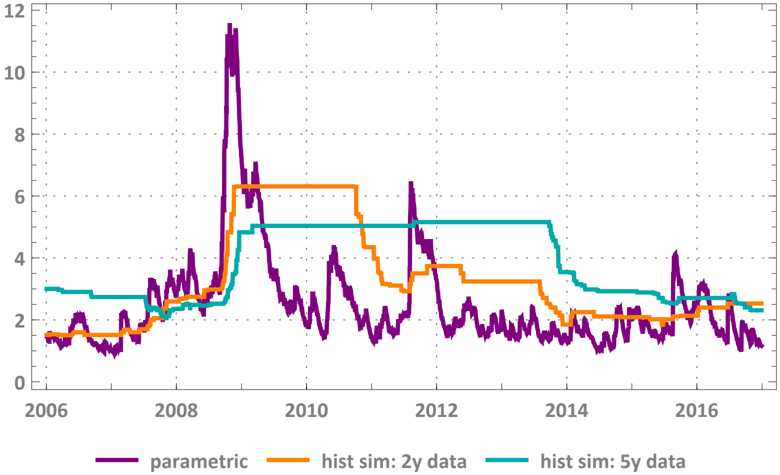

Time series of VaR estimates for long position in the S&P 500 index, daily, 03Jan2006 to 30Dec2016, expressed as returns in percent. Parametric estimates use a decay factor of 0.94, historical simulation estimates use 2 or 5 years of daily return data.

Dependence of VaR on confidence level and horizon

- Parametric VaR increases with both horizon τ and confidence level α
  - $z_{1-\alpha} = \Phi^{-1}(1-\alpha)$ becomes a larger-magnitude negative number as α increases
  - $-z_{1-\alpha}\sigma_t\sqrt{\tau}$ follows the square-root-of-time rule
- VaR computed via Monte Carlo and historical simulation increases with confidence level
- But if there is strong mean reversion in return volatility, VaR computed via historical simulation may be smaller at a longer than at a shorter horizon
- Table displays $-z_{1-\alpha}\sigma_t\sqrt{\tau}$ with $\sigma_t = 0.0069105$, for different values of τ and confidence level α (in percent)

	$\alpha = 0.95$	$\alpha = 0.99$	$\alpha = 0.995$
$\tau = 1$	1.13667	1.60762	1.78002
$\tau = 5$	2.54168	3.59475	3.98025