finvalrisk — Value at risk+

+This command is part of StataNow.

Description

`finvalrisk` calculates value at risk (VaR) using the historical, normal-based, and model-implied methods.

Quick start

Compute historical and normal-based VaR for assets `p1` through `p4`

```
finvalrisk p1-p4
```

Same as above, but compute risk at the 1st and 5th percentiles

```
finvalrisk p1-p4, percentiles(1 5)
```

Use confidence level instead of percentile; calculate VaR at 90%, 95%, and 99% confidence levels

```
finvalrisk p1-p4, levels(90 95 99)
```

Compute ARCH-based VaR

```
finvalrisk p1-p4, model(arch , ar(1) arch(1) garch(1))
```

Menu

Statistics > Financial statistics > Value at risk

Syntax

finvalrisk *varlist* [*if*] [*in*] [, *options*]

options	Description
Main	
percentiles(*numlist*)	compute VaR for specified percentiles; default is percentiles(5)
levels(*numlist*)	compute VaR for specified confidence levels
altdef	use alternative percentiles formula for historical estimate
model(*model_spec*)	specify conditional mean model or conditional variance model or both
novarlabel	display variable names rather than variable labels

You must tsset your data before using finvalrisk; see [TS] **tsset**.

collect is allowed; see [U] **11.1.10 Prefix commands**.

model_spec	Description
arima [, *arima_options*]	fit an autoregressive integrated moving-average (ARIMA) model
ewma [#] [, noconstant]	fit an exponentially weighted moving-average (EWMA) model
arch [, *arch_options*]	fit an autoregressive conditional heteroskedasticity (ARCH) model

Options

Main

percentiles(*numlist*) specifies the percentiles at which VaR should be computed. Specifying percentiles(1) gives the first (1%) percentile. Specifying percentiles(0.1) gives the 0.1% percentile. The default is percentiles(5). percentiles() may not be specified with levels().

levels(*numlist*) specifies the confidence levels at which VaR should be computed. levels(90) specifies that the VaR should be computed using a 90% confidence level; this is equivalent to specifying percentiles(10). levels() may not be specified with percentiles().

altdef uses an alternative formula for calculating percentiles when the historical estimate is requested. The default method is to invert the empirical distribution function by using averages, $(x_i + x_{i+1})/2$, where the function is flat (the default is the same method used by summarize; see [R] **summarize**). The alternative formula uses an interpolation method. See *Methods and formulas* in [D] **pctile**.

model(*model_spec*) calculates VaR based on a model. Three models are supported: arima, ewma, and arch.

model(arima [, *arima_options*]) specifies that returns follow an ARIMA process. All options allowed with the arima command are allowed; see [TS] **arima**.

model(ewma [#] [, noconstant]) specifies that the variance follows an EWMA model, $\sigma_t^2 = \omega + \lambda\sigma_{t-1}^2 + (1-\lambda)e_{t-1}^2$, where e_t is the error term. The parameter λ can be set to #. If it is not set, it is estimated. noconstant suppresses the parameter ω; otherwise, it is estimated.

model(arch [, *arch_options*]) specifies that returns follow an ARCH process. All options allowed with the arch command are allowed; see [TS] **arch**. Of particular interest are the three distributions of the error process available for arch: normal, Student's t, and the generalized error distribution.

novarlabel specifies that the variable names be displayed instead of the variable labels. By default, variable labels are used to label assets in the output.

Remarks and examples

Remarks are presented under the following headings:

Introduction

In financial analysis, it is often just as important to quantify potential losses as it is to summarize average behavior. Risk management methods seek to characterize low-probability, highly unfavorable events. The notion of VaR provides one measure of risk that is commonly used in practice.

VaR is a measure of potential loss on the value of an asset or a portfolio over a specified length of time for a given confidence level. It is often expressed as a number representing an amount of loss (or percentage loss) at a particular confidence level for a particular period. If the VaR on a portfolio is $100 at a 95% confidence level at a one-month horizon, then there is a 5% chance of losing $100 or more by the end of the month. The same portfolio might have a $300 VaR at the 99% confidence level at one month, meaning we expect a 1% chance of losing $300 or more by the end of the month. Put another way, losses more extreme than the specified VaR can still occur but should occur with less than the chosen percentile probability over the desired time frame.

The term "risk" can encompass multiple types of measurable uncertainty. Because VaR uses the observed historical performance of an asset or portfolio to assess risk, it is sometimes said to capture "market" risk and not, for example, counterpart, operational, or model risk (Boffelli and Urga 2016, chap. 5).

VaR can be thought of as an estimate of one or more quantiles of the distribution of asset returns. To estimate the quantiles, a model of potential returns on the asset or portfolio is required. Historical returns allow the model parameters to be fit to data. Many models have been proposed. The simplest is to use historical percentiles of the return distribution. Another approach is to estimate the parameters of a normal distribution based on historical returns and use the normal quantiles to compute VaR. This method, however, assumes time-invariant volatility, as well as normality of returns, neither of which may hold. Model-based approaches can be used to mitigate against the constant conditional mean assumption and the constant conditional variance assumption.

`finvalrisk` implements five models for measuring VaR:

1. Historical quantiles.
2. Normal quantiles based on the parameters estimated from historical data.
3. Quantiles derived from an ARIMA model for the mean of returns and a normally distributed error structure.
4. Quantiles derived from an EWMA model for the conditional variance.

5. Quantiles derived from an ARCH model for the conditional mean and conditional variance, with a choice of three error families for the error term: normal, Student's t, and the generalized error distribution.

We stress that all of these procedures rely on the assumption that past returns are a useful basis for forecasting the distribution of future returns.

Historical and normal-based VaR

The historical method of VaR uses observed quantiles of the distribution of returns. The normal-based approximation uses the mean and variance of the distribution of returns to compute a normal quantile.

We will compute the historical VaR and normal-based VaR for some fictional assets. The object of interest is the monthly percentage return on the asset, so the VaR will be a number like 1.5% at the 95% confidence level, meaning that there is a 5% chance of a loss of 1.5% or more. This VaR is sometimes referred to as a 5% VaR (based on the percentile) and sometimes as a 95% VaR (based on the confidence level). We follow Hurn et al. (2020) and call it a 5% VaR.

We can convert the VaR figures into dollar amounts by simply multiplying the percentage loss by the value of the asset.

Example 1: Historical return distribution

This example explores VaR graphically and numerically. We import some fictional stock prices. The data consist of the end-of-month price for 25 fictional publicly traded firms and range from 1955 through 2019. We use `finreturns` to generate monthly returns.

```
. use https://www.stata-press.com/data/r19/finex
(Fictional stock price data)
. quietly finreturns acme-tks, simple(r_) multiply(100)
```

`finreturns` created a set of variables `r_acme`, `r_bat`, etc, one for each asset. These variables contain the monthly simple return, multiplied by 100. So a value of 1.5 indicates that the asset grew in value by 1.5% in that month.

Somewhat arbitrarily, we choose asset `r_aaa` and observe its characteristics.

```
. summarize r_aaa, detail
                            AAA
-------------------------------------------------------------
      Percentiles      Smallest
 1%    -5.246088      -10.97514
 5%    -3.031731      -8.011154
10%    -1.930579      -6.801289       Obs                 779
25%    -.5408238      -6.047291       Sum of wgt.         779

50%     .7043946                      Mean           .5583062
                        Largest       Std. dev.       2.04691
75%     1.851859       5.240069
90%     2.941336       5.536678       Variance       4.189842
95%      3.61595       6.112684       Skewness      -.6100604
99%     4.996573       7.364605       Kurtosis       4.899915
```

The mean return for this asset is 0.56% per month, about one half of one percent per month, or about 6.9% per year, annualized using $1.0056^{12} - 1$. The percentiles give empirical quantiles of the distribution. We see that 1% of the time, this stock loses 5.25% or more of its value in a month. The percentiles reported in the table are negative, representing losses, but the VaR is typically reported as a positive value. Thus, the asset's 1% VaR is 5.25%. Similarly, 5% of the time, the asset loses 3.03% or more in a month. Its 5% VaR is 3.03%.

The same information can be conveyed graphically, with a vertical line marking the 5% quantile of the distribution:

```
. histogram r_aaa, title("Histogram of returns") xline(-3.032)
> xlabel(-10 -5 -3.03 0 5 10) color(stblue%20)
(bin=27, start=-10.975141, width=.67924986)
```

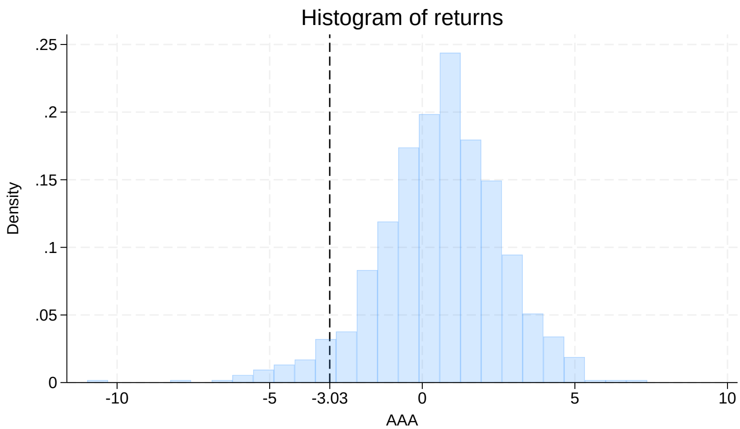

The observed returns are on average positive, though with a substantial fraction of months with negative return. The 5% VaR line marks the point at which 5% of returns are at that level of return or lower. We see there is a somewhat long left tail, reminding us that the 5% VaR is not the minimum return: 5% of returns are worse than it and indeed can be much worse.

Example 2: Normal-based VaR

Continuing with example 1, we see the normal-based approximation of VaR fits a normal distribution to the observed mean and variance of the returns:

```
. histogram r_aaa, title("Histogram of returns") normal color(stblue%20)
(bin=27, start=-10.975141, width=.67924986)
```

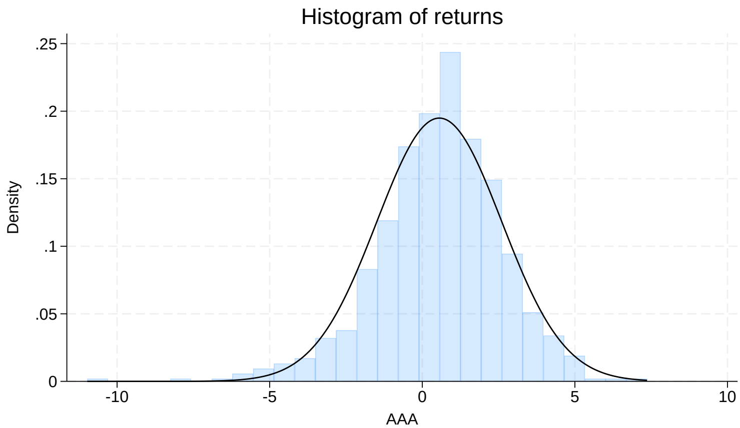

We could now compute the implied 5% normal quantile. Instead, we ask `finvalrisk` to do it for us. `finvalrisk` takes a variable list, which contains only `r_aaa` in this case. By default, the 5% VaR is reported, which corresponds to the 95% confidence level.

```
. finvalrisk r_aaa

Value-at-risk percentiles

Number of obs = 779
Sample: 1955m2 thru 2019m12
```

	Empirical	Normal
5%		
AAA	-3.031731	-2.808562

Both the historical and normal-based quantiles are reported. The normal-based VaR is 2.8%, a somewhat less extreme estimate than the empirical distribution suggests. Because our asset has a long left tail, the normal approximation underestimates VaR relative to the percentile method.

Example 3: Specifying multiple assets and thresholds

Continuing with the previous example, we now explore VaR for several assets and several thresholds. We extend the exploration to include a simple portfolio of assets. Using `finportfolio`, we can take an equally weighted average of four of the assets and call it `p1`.

```
. finportfolio equal r_aaa-r_cph, generate(p1, label("Portfolio"))
Equally weighted portfolio

Number of obs = 779
Sample: 1955m2 thru 2019m12
```

	Weight
AAA	.25
AFH	.25
ARD	.25
CPH	.25

```
Portfolio return    = 0.4809
Portfolio std. dev. = 1.6167
Risk-free rate      = 0.0000
Sharpe ratio        = 0.2975
```

Now we specify the `percentiles()` option to compute the 5% and 1% VaRs (or, equivalently, VaRs at the 95% and 99% confidence levels) for the individual assets and the portfolio.

```
. finvalrisk r_aaa-r_cph p1, percentiles(5 1)
Value-at-risk percentiles

Number of obs = 779
Sample: 1955m2 thru 2019m12
```

	Empirical	Normal
5%		
AAA	-3.031731	-2.808562
AFH	-3.629369	-3.287768
ARD	-2.758431	-2.796273
CPH	-2.991949	-3.005515
Portfolio	-2.366248	-2.178334
1%		
AAA	-5.246088	-4.203519
AFH	-5.519961	-4.923469
ARD	-4.401581	-4.071673
CPH	-5.37588	-4.426132
Portfolio	-4.672417	-3.280123

The output table has two blocks, one for each VaR threshold. At the 5% level, the four individual assets have VaRs ranging from 3.63% to 2.76%. The equally weighted portfolio achieves a less extreme VaR, 2.37%. However, portfolios are not guaranteed to beat individual assets, especially at the tails. At the 1% VaR threshold, the portfolio has a VaR of 4.67%, but one of the assets (ARD) has a slightly lower VaR of 4.40%.

Model-based VaR

So far, we have used the historical distribution of returns and a normal-based approximation to that distribution to compute VaR. A more sophisticated approach involves estimating the conditional mean or conditional variance of returns, or both, as time-varying objects. Doing so generates a VaR that can fluctuate over time as the behavior of returns changes. Three models are available. The ARIMA model is a model for the conditional mean. Intuitively, if stock returns have any predictability on average, then the expected mean return can shift over time. If the expected mean is unusually high or low in a given period, it will shift the whole distribution and thus shift VaR. The EWMA model and ARCH model are models for the conditional variance. Intuitively, many stock returns exhibit clusters of high or low volatility. In high-volatility periods, the distribution of returns is itself wider, thus increasing VaR; in low-volatility periods, the distribution of returns is narrower, reducing VaR.

Example 4: ARIMA for the mean

In this example, we continue with example 3 to explore VaR for the four sample assets and their equally weighted portfolio. We begin by asking whether a model for the average return adds any nuance to the VaR calculations. We consider a simple model for the mean where returns r_{it}'s are modeled based on their first and second lags and where errors e_{it}'s are assumed to be normally distributed.

$$r_{it} = \beta_0 + \beta_1 r_{i,t-1} + \beta_2 r_{i,t-2} + e_{it}$$
$$e_{it} \sim N(0, \sigma^2)$$

This is a simple AR(2) model, and its parameters β_0, β_1, β_2, and σ^2 can be estimated with `arima`. We fit this model to one of the assets, `r_aaa`.

```
. arima r_aaa, ar(1/2) nolog

ARIMA regression

Sample: 1955m2 thru 2019m12                     Number of obs     =        779
                                                Wald chi2(2)      =      34.22
Log likelihood = -1647.842                      Prob > chi2       =     0.0000

------------------------------------------------------------------------------
             |                 OPG
       r_aaa | Coefficient  std. err.      z    P>|z|     [95% conf. interval]
-------------+----------------------------------------------------------------
r_aaa        |
       _cons |   .5579395   .0869368     6.42   0.000     .3875466    .7283325
-------------+----------------------------------------------------------------
ARMA         |
          ar |
         L1. |   .1951391   .0337801     5.78   0.000     .1289313    .2613469
         L2. |  -.0676834   .0351207    -1.93   0.054    -.1365187     .001152
-------------+----------------------------------------------------------------
      /sigma |   2.006473   .0415384    48.30   0.000     1.925059    2.087887
------------------------------------------------------------------------------
Note: The test of the variance against zero is one sided, and the two-sided
      confidence interval is truncated at zero.
```

There is some mild forecastability: the first autoregressive lag coefficient is about 0.2, so when returns are 1% higher than their long-run average in a given period, then returns are expected to be 0.2% above their long-run mean in the following period.

To use `finvalrisk` to fit this AR(2) model to each asset and the portfolio, we specify the `model()` option, select a model, and provide any model-specific options. We obtain the following output:

```
. finvalrisk r_aaa-r_cph p1, model(arima, ar(1/2))
Fitting ARIMA models ...
Value-at-risk percentiles
Model: ARIMA
Error family: Normal
Number of obs = 779
Sample: 1955m2 thru 2019m12
```

	Empirical	Normal	ARIMA historical	ARIMA one step
5%				
AAA	-3.031731	-2.808562	-2.742248	-2.922465
AFH	-3.629369	-3.287768	-3.263831	-3.238124
ARD	-2.758431	-2.796273	-2.791969	-2.763895
CPH	-2.991949	-3.005515	-2.964135	-2.702158
Portfolio	-2.366248	-2.178334	-2.106934	-2.028872

Note: One-step-ahead value at risk for **2020m1**.

The output now has four columns, one for each VaR method. As before, the first column reports the historical quantile, and the second column reports the normal-based quantile.

The third column reports the quantile for the model-based VaR averaged across the sample. The final column provides the quantile for the one-step-ahead VaR forecast for one period out from the sample period. For example, these fictional data end in December 2019, so the one-step-ahead quantile is reported for January 2020.

Example 5: EWMA for the variance

The EWMA model allows a time-varying error variance. Specifically, instead of assuming that the variance is constant at σ^2, the error variance is allowed to depend on past error variance and the most recent error value:

$$\sigma_t^2 = \omega + \lambda\sigma_{t-1}^2 + (1-\lambda)e_{t-1}^2$$

The time-varying variance allows for periods of high volatility and low volatility. In a period of high volatility, the distribution of returns is more spread out, and VaR will be more extreme than average. In periods of calm, the distribution of returns is less spread out, and VaR will be less extreme than average.

`finvalrisk` implements this method in the `model(ewma [#])` option, where the optional number is the weight (or shape parameter) λ. If a weight is not given, it is calculated optimally from the data.

```
. finvalrisk r_aaa-r_cph p1, model(ewma 0.95)

Fitting EWMA models ...

Value-at-risk percentiles

Model: EWMA
Error family: Normal
Shape parameter = 0.95

Number of obs = 779
Sample: 1955m2 thru 2019m12
```

	Empirical	Normal	EWMA historical	EWMA one step
5%				
AAA	-3.031731	-2.808562	-2.92945	-2.431956
AFH	-3.629369	-3.287768	-3.35059	-2.619925
ARD	-2.758431	-2.796273	-2.905166	-2.423582
CPH	-2.991949	-3.005515	-3.142034	-3.021991
Portfolio	-2.366248	-2.178334	-2.277888	-1.991785

Note: One-step-ahead value at risk for **2020m1**.

The header reports the EWMA shape parameter of 0.95, which we specified. In the output table, the third column gives the quantile corresponding to the historical average VaR from the EWMA model for each asset. The fourth column gives the one-step-ahead value for the period immediately following the last sample period.

The portfolio has an empirical historical VaR of 2.37%. The average EWMA estimated VaR is a little lower, 2.28%. But the one-step-ahead forecasted VaR for January 2020 is 2.0%, indicating that, given the most recent observations of 2019, January 2020 is predicted to be a quieter month.

Example 6: ARCH for the mean and variance

In this example, we fit an ARCH model of the conditional mean and conditional variance. We first do so by hand and compute time-varying VaR by hand to provide intuition. We fit a generalized ARCH(1, 1) model with the `arch` command, using the `arch(1)` and `garch(1)` options.

```
. arch r_aaa, ar(1) arch(1) garch(1) nolog
ARCH family regression -- AR disturbances
Sample: 1955m2 thru 2019m12                     Number of obs     =        779
                                                Wald chi2(1)      =      13.22
Log likelihood = -1636.907                      Prob > chi2       =     0.0003
```

r_aaa	Coefficient	OPG std. err.	z	P>|z|	[95% conf.	interval]
r_aaa						
_cons	.5896113	.0860708	6.85	0.000	.4209157	.758307
ARMA						
ar L1.	.1525666	.041967	3.64	0.000	.0703129	.2348204
ARCH						
arch L1.	.0883585	.0241089	3.66	0.000	.0411059	.1356112
garch L1.	.8003264	.0692668	11.55	0.000	.664566	.9360868
_cons	.4555743	.2151189	2.12	0.034	.0339489	.8771997

The ARCH and GARCH terms are both positive, indicating that the variance of the series exhibits some positive autocorrelation, so periods of high and low variance are clustered together.

To compute VaR by hand, we use `predict` after `arch` to predict the conditional mean and variance. Then we build VaR using the normal error distribution assumption.

```
. predict raaa_mean, xb
. predict raaa_variance, variance
. generate raaa_VaR_arch = raaa_mean + invnormal(0.05) * sqrt(raaa_variance)
```

These three steps compute percentiles corresponding to the 5% ARCH-based VaR. We can now graph it, along with the percentile corresponding to historical 5% VaR as a benchmark. We first generate a variable with the 5th percentile for the historical VaR, and then we add some variable labels:

```
. label variable raaa_VaR_arch "5th percentile, ARCH VaR"
. generate raaa_VaR_hist = -3.03
. label variable raaa_VaR_hist "5th percentile, historical VaR"
```

The following graph plots returns over time, along with loss corresponding with ARCH and historical VaR.

```
. tsline r_aaa raaa_VaR_hist raaa_VaR_arch, legend(rows(1))
> title("Returns and VaR percentile estimates")
```

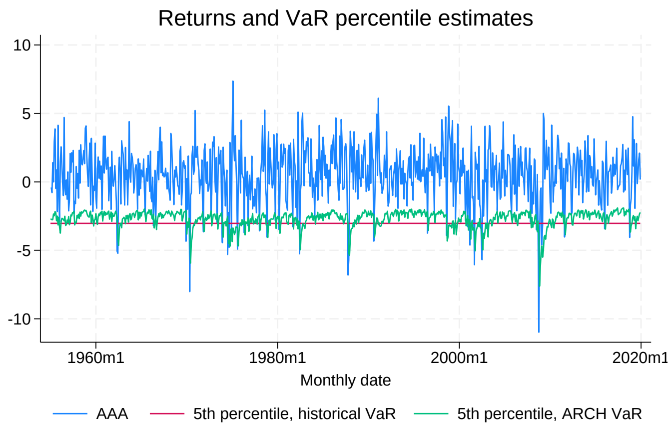

We see that the volatility of the series itself varies over time, with some periods of relatively high volatility and other periods of relative calm. The potential loss based on the historical VaR is a straight line at the 5% quantile of the distribution. The potential loss based on the ARCH-based VaR adapts over time to the recent conditions of the series. For example, when returns are especially volatile, potential loss becomes more negative, reflecting higher risk in those time periods. We next use `finvalrisk` to perform these calculations.

```
. finvalrisk r_aaa-r_cph p1, model(arch, ar(1) arch(1) garch(1))
Fitting ARCH models ...
Value-at-risk percentiles
Model: ARCH
Error family: Normal
Number of obs = 779
Sample: 1955m2 thru 2019m12
```

	Empirical	Normal	ARCH historical	ARCH one step
5%				
AAA	-3.031731	-2.808562	-2.701858	-2.433388
AFH	-3.629369	-3.287768	-3.166423	-2.750784
ARD	-2.758431	-2.796273	-2.773078	-2.344072
CPH	-2.991949	-3.005515	-2.923198	-2.515059
Portfolio	-2.366248	-2.178334	-2.048298	-1.665521

Note: One-step-ahead value at risk for **2020m1**.

The historical and normal-based VaR are the same as we have already encountered. The 5% quantile for the ARCH-based VaR for the historical sample and for the one-step-ahead prediction are reported in the third and fourth columns, respectively. Notice that while the historical ARCH average VaR for the portfolio is 2.05%, the one-step-ahead conditional VaR is much smaller at 1.67%, indicating that the upcoming period is predicted to be less volatile than the historical average. This prediction would be driven by the underlying ARCH model.

Alternative error distributions

Although we have so far allowed time-varying variance, the underlying model for the error terms remains a normal distribution. The `arch` command allows three distributions: normal, Student's t, and the generalized error distribution. The last two both have shape parameters, which can be controlled manually or estimated. The shape parameter of the t distribution is its degrees of freedom. The shape parameter for the generalized error distribution controls the thickness of the tails. When the shape parameter of the generalized error distribution is less than two, its tails are fatter than a normal distribution; when the shape parameter is greater than two, the tails are thinner than the normal distribution.

Example 7: ARCH with GED errors

To specify the t distribution, we would use `distribution(t)`. But we are interested in the generalized error distribution, so we type `distribution(ged)`. We also specify `percentiles(1)` because we are interested in the 1% VaR.

```
. finvalrisk r_aaa-r_cph p1,
> model(arch, ar(1) arch(1) garch(1) distribution(ged)) percentiles(1)
Fitting ARCH models ...
Value-at-risk percentiles
Model: ARCH
Error family: Generalized error
Shape parameter = 1.60
Number of obs = 779
Sample: 1955m2 thru 2019m12
```

	Empirical	Normal	ARCH historical	ARCH one step
1%				
AAA	-5.246088	-4.203519	-4.344637	-3.953778
AFH	-5.519961	-4.923469	-5.097832	-4.542212
ARD	-4.401581	-4.071673	-4.054237	-3.442605
CPH	-5.37588	-4.426132	-4.57175	-4.116522
Portfolio	-4.672417	-3.280123	-3.434237	-3.003905

Note: One-step-ahead value at risk for **2020m1**.

In the output header, Stata reports we are using the generalized error distribution. The shape parameter is estimated separately for each return or portfolio, and the reported shape parameter is the average of the estimated shape parameters. The average here is 1.60, which indicates that the returns and portfolios considered here are estimated to have fatter tails than the normal distribution would predict. VaR calculations, both historical average and one-period-ahead, account for these fatter tails. For example, the historical average ARCH VaR for asset AAA is 4.34%, whereas the normal distribution approach has VaR of 4.20%.

You may wish to set the shape parameter instead of estimating it. To do so, you can ask for an ARCH model with a specific distribution(ged #).

Example 8: Fixing the shape parameter

We specify a relatively small value of the shape parameter, 1.20, to see how it affects VaR estimation. Unlike the previous example, the shape parameter is fixed and fixed to the same value for all assets.

```
. finvalrisk r_aaa-r_cph p1,
> model(arch, ar(1) arch(1) garch(1) distribution(ged 1.2)) percentiles(1)

Fitting ARCH models ...

Value-at-risk percentiles

Model: ARCH
Error family: Generalized error
Shape parameter = 1.20

Number of obs = 779
Sample: 1955m2 thru 2019m12
```

	Empirical	Normal	ARCH historical	ARCH one step
1%				
AAA	-5.246088	-4.203519	-4.794479	-4.403476
AFH	-5.519961	-4.923469	-5.76262	-5.182215
ARD	-4.401581	-4.071673	-4.901967	-4.184606
CPH	-5.37588	-4.426132	-5.152057	-4.695892
Portfolio	-4.672417	-3.280123	-3.63922	-3.222156

```
Note: One-step-ahead value at risk for 2020m1.
```

As expected, specifying an error distribution with fatter tails leads to an increased estimated downside risk, and VaR for the ARCH models grows. For instance, a 1% VaR is now associated with a 4.80% decline in the AAA asset returns on average, compared with the normal distribution estimate of 4.20%.

Stored results

finvalrisk stores the following in r():

Scalars

r(N)	number of observations
r(tmin)	minimum time
r(tmax)	maximum time
r(shape)	shape parameter for models arch and ewma
r(altdef)	1 if altdef is specified, 0 otherwise

Macros

r(model)	arch or arima or ewma
r(modelopts)	model-specific options
r(tmins)	formatted minimum time
r(tmaxs)	formatted maximum time
r(percentiles)	requested percentiles
r(levels)	requested confidence levels

Matrices

r(table)	matrix containing empirical, normal, and model-based value-at-risk statistics

Methods and formulas

Methods and formulas are presented under the following headings:

Historical quantiles

`finvalrisk` uses the `_pctile` command to estimate historical quantiles. The default formula for percentiles is as follows. For each series x_t, let $x_{(j)}$ refer to the x in ascending order for $j = 1, 2, \ldots, T$. Hence, $x_{(j)}$ is the jth largest x. To obtain the pth percentile, which we will denote as $x_{|p|}$, let $P = Tp/100$. Let i be the integer floor of $Tp/100$. The pth percentile is then

$$x_{[p]} = \begin{cases} \dfrac{x_{(i)} + x_{(i+1)}}{2} & \text{if } i = P \\ x_{(i+1)} & \text{otherwise} \end{cases}$$

When the `altdef` option is specified, the following alternative definition is used. Let i be the integer floor of $(T+1)p/100$; that is, i is the largest integer $i \leq (T+1)p/100$. Let h be the remainder $h = (T+1)p/100 - i$. The pth percentile is then

$$x_{[p]} = (1-h)x_{(i)} + hx_{(i+1)}$$

where $x_{(0)}$ is taken to be $x_{(1)}$ and $x_{(T+1)}$ is taken to be $x_{(T)}$. See [D] **pctile** for further information on `_pctile`.

The underlying modeling assumption is that the historical distribution of returns is stable and that the next period's return is a random draw from the historical distribution. Thus, the p% VaR is the pth quantile of the historical distribution.

Normal quantiles

Estimates of the normal quantile are a function of the estimated mean, the estimated standard deviation, and a normal density. The formula for the pth percentile for variable x_t is

$$x_p = \hat{\mu} + \hat{\sigma}\Phi^{-1}(p/100)$$

where $\hat{\mu}$ is the estimated mean of x and $\hat{\sigma}$ is the estimated standard deviation of x. The function $\Phi^{-1}(p)$ is the inverse cumulative standard normal distribution; if $\Phi(z) = p$, then $\Phi^{-1}(p) = z$.

The underlying modeling assumption is that returns are independent, identically distributed draws from a normal distribution. Once the mean μ and standard deviation σ have been estimated, the p% VaR is the pth quantile of the normal distribution. This method may understate VaR if the actual distribution is nonnormal or if the mean or variance of the distribution changes over time.

Model-based quantiles

Model-based quantiles use a model to obtain potentially time-varying estimates of the mean return $\hat{\mu}_t$ and the standard deviation $\hat{\sigma}_t$, then combine these estimates with a distribution for the model errors. Because the conditional mean and conditional standard deviation can be time varying, the quantiles of the expected return distribution, and thus VaR, can also be time varying.

Regardless of the model used, `finvalrisk` reports two numbers: the average and one-step-ahead VaR. The average model-based VaR is the average of the period-by-period VaR estimates. The one-step-ahead VaR is the VaR predicted for one period after the sample ends, using the one-step out-of-sample estimate of the conditional mean and conditional standard deviation.

ARIMA

`finvalrisk` with the `model(arima , ...)` option fits an ARIMA model for each series x_t and uses the model to obtain quantile estimates.

The conditional mean $\hat{\mu}_t$ and conditional standard deviation $\hat{\sigma}_t$ are predicted each period using the fitted model. Then a value-at-risk estimate for each period is constructed from the estimated mean, estimated standard deviation, and the normal distribution.

ARCH

`finvalrisk` with the `model(arch, ...)` option estimates an ARCH model for each series x_t and uses the model to form quantile estimates.

One of the simpler ARCH models is the generalized ARCH(1,1), in which the conditional variance evolves over time via

$$\sigma_t^2 = \omega + \alpha e_{t-1}^2 + \beta \sigma_{t-1}^2$$

where (ω, α, β) are parameters to be estimated and e_t^2 is the squared error from the model. Many more complicated models can be fit; see [TS] **arch**.

The conditional mean $\hat{\mu}_t$ and conditional standard deviation $\hat{\sigma}_t$ are predicted each period using the fitted model. Then a value-at-risk estimate for each period is constructed from the estimated mean, estimated standard deviation, and one of three distributions. The average model-based quantile is the average of the time series of VaR estimates. The available distributions are a normal distribution, a t distribution, and a generalized error distribution. The t distribution and generalized error distribution each depend on a shape parameter. This parameter can be fixed or estimated.

EWMA

The EWMA process for the variance is

$$\sigma_t^2 = \omega + \lambda \sigma_{t-1}^2 + (1 - \lambda) e_{t-1}^2$$

The parameter λ can be fixed or estimated. The parameter ω can be estimated or fixed to 0 with the `noconstant` option. Then quantiles are computed from a normal distribution with mean $\hat{\mu}$ and standard deviation $\hat{\sigma}_t$.

References

Boffelli, S., and G. Urga. 2016. *Financial Econometrics Using Stata*. College Station, TX: Stata Press.

Hurn, S., V. L. Martin, P. C. B. Phillips, and J. Yu. 2020. *Financial Econometric Modeling*. Oxford: Oxford University Press.

Also see

[FIN] **finportfolio** — Financial portfolio selection+

[FIN] **finreturns** — Generate financial returns+