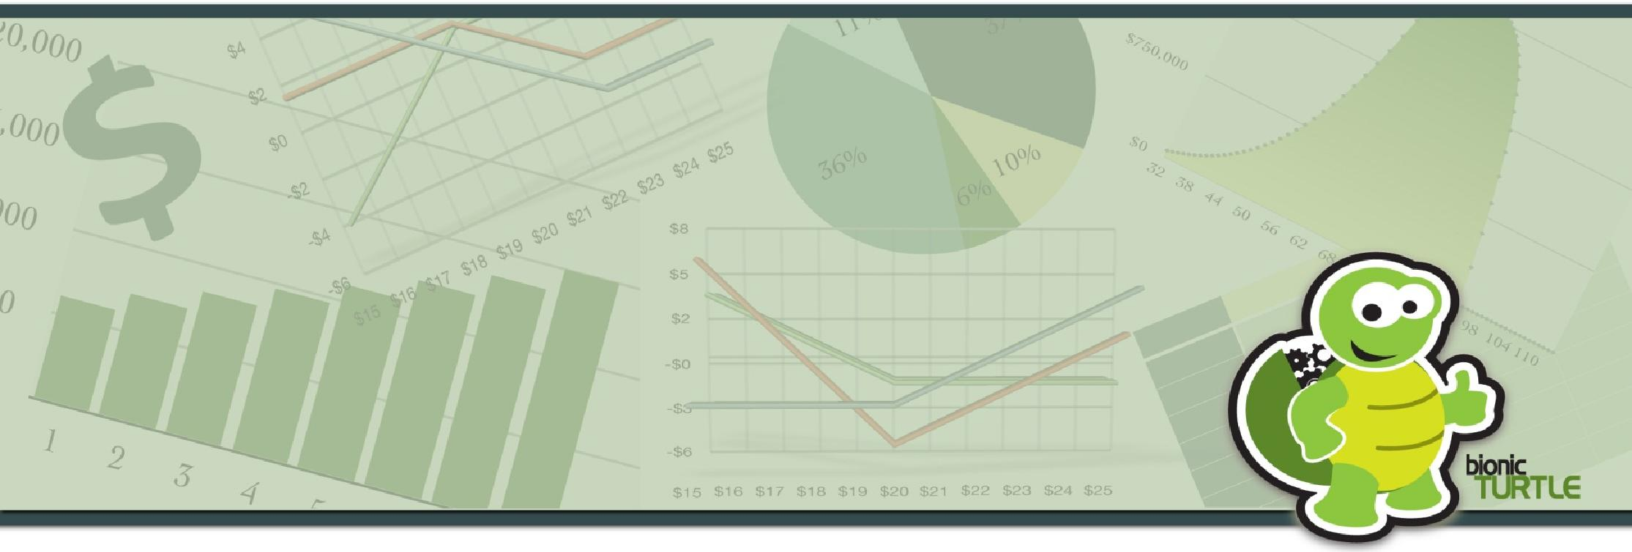

P1.T4. Valuation & Risk Models

Bionic Turtle FRM Practice Questions

Chapter 3. Measuring and Monitoring Volatility

By David Harper, CFA FRM CIPM
www.bionicturtle.com

Chapter 3. Measuring and Monitoring Volatility

Chapter 3. Measuring and Monitoring Volatility

P1.T4.803. Quantifying volatility in value at risk (VaR) models

Learning objectives: Explain how asset return distributions tend to deviate from the normal distribution. Explain reasons for fat tails in a return distribution and describe their implications. Distinguish between conditional and unconditional distributions. Describe the implications of regime switching on quantifying volatility.

803.1. Peter the Risk Analyst has collected one year (N = 250 days) of asset price data. He translates the daily prices into daily log returns; e.g., if a yesterday's price close was $5.00 and today's price close is $4.90, then the daily log return is LN(4.90/5.00) = -2.020%. His initial exploratory data analysis (EDA) of this return series includes the following observations:

- The lag-1 autocovariance is approximately zero
- The kurtosis of the returns distribution is about 9.50
- The daily sample standard deviation of returns is about 1.50%

In this situation, each of the following statements is true or at least plausible **EXCEPT** which statement is inaccurate?

a) The unconditional distribution obviously is not normal, but the unconditional distribution might fit a normal mixture distribution
b) A plausible explanation for the observed unconditional heavy tails is a distribution that is conditionally normal yet time-varying or regime-shifting
c) The zero autocovariance supports (i.e., is necessary but not sufficient to) scaling the daily standard deviation to a per annum volatility of about 23.72% = 1.50% * SQRT(250)
d) If he imposes normality by standardizing the returns per (X - μ)/σ, then Peter can expect to lower kurtosis of the returns distribution to approximately three (~3.0) and this will also shift more probability mass to the shoulders of the distribution

803.2. The daily standard deviation of a risky asset is 1.40%. If daily returns are independent, then of course we can use the square root rule (SRR) to scale into a two-day volatility given by 1.40%*SQRT(2) = 1.980%. However, if the lag-1 autocorrelation of returns is significantly positive at 0.60, then which is **NEAREST** to the corresponding scaled two-day volatility?

a) a. 1.64%
b) b. 1.98%
c) c. 2.33%
d) d. 2.50%

803.3. Based on daily asset prices over the last month (n = 20 days), Rebecca the Risk Analyst has calculated the daily volatility as a basic historical standard deviation (abbreviated STDEV to match Linda Allen's label for the equally weighted average of squared deviations which "is the simplest and most common way to estimate or predict future volatility"). As the exhibit below shows, however, she could not quite decide which mean, μ, to use for the sample standard deviation: the actual sample mean return was -45.12 basis points; but she has also seen zero assumed per squared returns; and finally, the expected return is actually +20.0 bps. Each produces a slightly different sample standard deviation. For example, assuming the actual sample mean produces a daily volatility of 56.04 bps which matches the Excel function STEV.S().

				Mean, μ, assumption used in computed volatility		
				Sample	Zero	Expected
		Return		-45.12	0.00	20.00
	Price	(bps)		(R- -45.1)^2	(R- 0.0)^2	(R- 20.0)^2
today	$9.50	-32.9		148.2	1,085.7	2,803.6
t - 1	$9.53	-14.7		927.5	215.1	1,201.7
t - 2	$9.55	13.0		3,380.5	169.5	48.7
t - 3	$9.53	-102.6		3,305.2	10,529.4	15,033.9
	...				...	
t - 17	$10.20	-26.6		344.7	705.1	2,167.3
t - 18	$10.23	-133.5		7,805.9	17,814.9	23,553.8
t - 19	$10.37	-31.0		198.3	963.5	2,605.1
t - 20	$10.40					
	Avg	(45.1)	Sum:	59,677.1	100,396.0	144,493.1
	Sample σ	56.04	÷ 19	3,140.90	5,284.00	7,604.90
			Sample σ	56.04	72.69	87.21

Which of the following statements is **TRUE** about Rebecca's calculations?

a) An advantage of this short 20-day window under the STDEV approach is that it minimizes the so-called ghosting effect
b) If the returns are especially heavy-tailed (ie, non-normal), this STDEV is a superior predictor to absolute standard deviation
c) If the actual 45.12 bps decline was uniquely unpredictable, it is reasonable to use the currently expected positive μ = +20 bps as the conditional mean parameter
d) While the assumption of zero mean is technically acceptable as an MLE (as opposed to unbiased) estimator of population standard deviation, it should in practice be avoided because we seek a conditional mean parameter

Answers:

803.1. D. False, and this is part of the theme of the early section in Linda Allen's Chapter 2. Standardizing simply re-scales the units so that a raw observation is expressed in standard normal units. This "imposes normality" but does not itself cause the data to become normal!

In regard to (A), (B) and (C) each is TRUE.

- ***In regard to true (A),*** the normal mixture distribution is very powerful and the mixture of two 5ormal will always generate a heavy-tailed distribution
- ***In regard to true (B),*** a conditional distribution that is normal but time-varying is an explanation consistent with Linda Allen's preferred interpretation of the observation of unconditionally heavy (non-normal) returns.
- ***In regard to true I,*** scaling per the square root rule (SRR) does not require normal assumption, but instead requires the independent and identically distributed (i.i.d.) returns. Zero autocovariance implies zero autocorrelation but does not imply independence (recall that independence → zero autocovariance but the converse is not necessarily true because dependence can be non-linear) such that zero autocovariance is necessary but not sufficient to scaling per the SRR. In practice, scaling thusly is commonly conducted.

803.2. D. 2.50%. The 2-period variance is given by σ^2[r(t,t+2)] = σ^2[r(t,t+1)] + σ^2[r(t+1,t+2) + 2*COV[r(t,t+1), r(t+1,t+2)]. In this case, 2-day volatility = SQRT(0.0140^2 + 0.0140^2 + 2*0.60*0.0140*0.0140) = SQRT(0.00062720) = 2.5044%

803.3. C. True: If the actual 45.12 bps decline was an uniquely unpredictable, it is reasonable to use the currently expected positive μ = +20 bps as the conditional mean parameter

In regard to (A), (B) and (D), each is FALSE:

- ***In regard to false (A),*** the STDEV approach with a short window maximizes the ghosting effect!
- ***In regard to false (B),*** writes Linda Allen, "it is worthwhile mentioning that an alternative procedure of calculating the volatility involves averaging absolute values of returns, rather than squared returns. This method is considered more robust when the distribution is non-normal. In fact it is possible to show that while under the normality assumption STDEV is optimal, when returns are non-normal, and, in particular, fat tailed, then the absolute squared deviation method may provide a superior forecast."
- ***In regard to false (D),*** zero mean is actually a very acceptable choice (and assuming zero mean does not imply an MLE estimator!)

Linda Allen: "Suppose, for example, that we need to estimate the volatility of the stock market, and we decide to use a window of the most recent 100 trading days. Suppose further that over the past 100 days the market has declined by 25 percent. This can be represented as an average decline of 25bp/day (−2,500bp/100days = −25bp/day). Recall that the econometrician is trying to estimate the conditional mean and volatility that were known to market participants during the period. Using −25bp/day as μt, the conditional mean, and then estimating $\sigma 2$ t, implicitly assumes that market participants knew of the decline, and that their conditional distribution was centered around minus 25bp/day.

Since we believe that the decline [in her example, the historical sample mean was entirely unpredictable, imposing our priors by using $\mu t = 0$ is a logical alternative. Another approach is to use the unconditional mean, or an expected change based on some other theory as the conditional mean parameter. In the case of equities, for instance, we may want to use the unconditional average return on equities using a longer period – for example 12 percent per annum, which is the sum of the average risk-free rate (approximately 6 percent) plus the average equity risk premium (6 percent). This translates into an average daily increase in equity prices of approximately 4.5bp/day. This is a relatively small number that tends to make little difference in application, but has a sound economic rationale underlying its use."

Discuss here in forum: https://www.bionicturtle.com/forum/threads/p1-t4-803-quantifying-volatility-in-value-at-risk-var-models-allen-ch-2.13519/