SOCIETY OF ACTUARIES

Article from:

Risk Management Newsletter

July 2004 – Issue No. 2

Getting to Know CTE

By David Ingram

The risk measure conditional tail expectation (CTE) has been getting more and more attention for measuring risk in any situation with non-normal distribution of losses. Canadian and U.S. insurance regulators have adopted CTE as a standard for regulatory capital measurement. Academics have lauded CTE as a "coherent" statistic. Those outside the insurance industry call it "Tail VaR" or "expected tail loss" (ETL). Actuaries, who have always been suspicious or even hostile to the usage of value at risk (VaR) as a risk measurement standard, have readily embraced CTE.

This article presents the observations of some random investigation into the nature of CTE. There are still many questions that need to be asked and answered before we can say that we are completely familiar with the nature and characteristics of CTE. There are six questions that will be discussed:

- How does CTE compare to VaR?
- How sensitive is CTE to the number of scenarios used?
- Is CTE consistent with existing RBC factors?
- What is the impact of time period used for calculating CTE?
- How does CTE measure the difference between regime switching and single regime lognormal scenarios?

These observations are all tied to examples. Since insurance products are very complex, it is difficult to see whether the results of any one test would apply to any other situation. Therefore, the examples used here will all be based on simple random values, rather than to calculations based on the random values. The random values are developed as stock returns that would be used as inputs to calculations about a stock-linked product. These discussions will not resolve any issues about CTE calculations for insurance products, but through some sort of rough transitive property, can form a part of a starting point for investigations into the ultimate nature of this new tool.

1. CTE vs. VaR

Without getting into the mathematics of coherence, a simple picture can show the appeal of CTE as opposed to VaR. CTE is defined as the probability weighted loss above a certain probability level, while VaR is the loss at a certain probability level. One definition of CTE is that it is the average of all VaR values for probabilities above a specified level. Generally VaR is used by banks and most often it is used to measure risk over a very short (less than one month) time frame. CTE is becoming the standard for insurance company risk measurement, especially regarding risk capital and is most often used to measure risk over multi-year time frames that are needed to view insurance risks. The following two charts show the distribution of values for stochastic tests of two products. Product A has a normal distribution of gains and losses, while the distribution of gains and losses for Product B are not as predictable.

Chart 1—Distribution of Gains and Losses

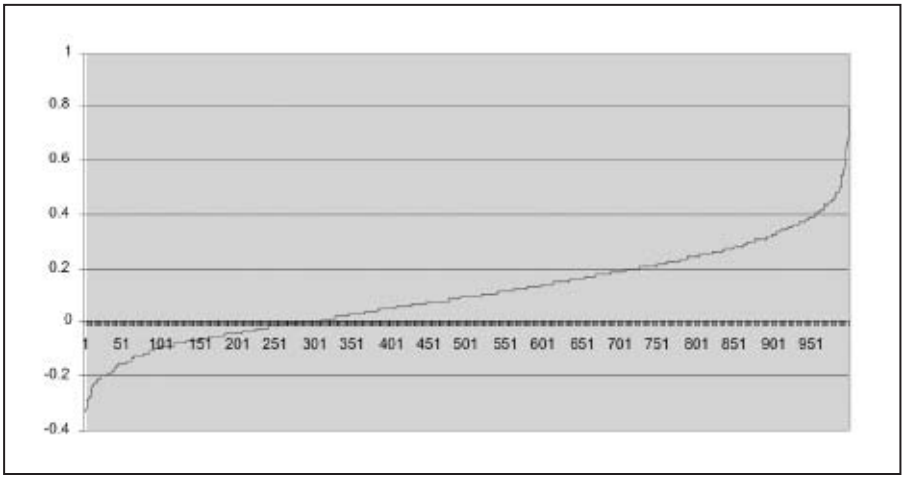

For Product A (Chart 1) and any other product with a normal distribution of gains and losses there is a regular predictable relationship between CTE and VaR. Using one measure or the other does not necessarily add any information.

continued on page 34 ◗

Chairperson David Ingram, FSA, MAAA, is a consulting actuary at Milliman USA in New York. He can be reached at david.ingram@milliman.com.

Getting to Know CTE

◗ continued from page 33

For Product B (Chart 2), VaR and CTE can sometimes show drastically different stories. The VaR for Product B on the graph above is a near zero value. There is a small loss compared to the expected gain, but not much. The tail of the profit distribution starts to drop-off somewhere above the 95th percentile. A CTE calculation includes the impact of that drop-off, wherever it starts. A VaR measure will only reflect the drop off if it is occurring at the percentage chosen for the VaR measure.

Above, we noted that there is a predictable relationship between VaR and CTE for a normally distributed gain/loss distribution. In fact, Chart 3 shows the values of 95VaR and CTE90 as a percentage of mean for various normal distributions are both linear relationships, with the 95VaR equal to the mean less 164.56 percent of the standard deviation and the CTE90 equal to the mean less 176.3 percent of the standard deviation. Looking at the graph in Chart 3, first notice that VaR and CTE are both positive when the standard deviation is less than about 65 percent of the mean. The modified CTE used for U.S. NAIC risk-based capital would use zero for any positive values and would therefore produce small negative values a little further to the left on the chart.

Chart 2—Embedded Value

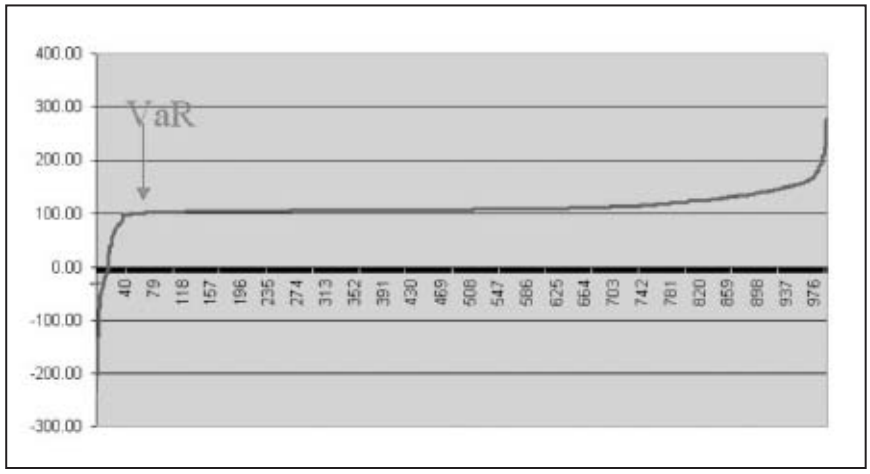

Chart 3—Normal Distribution

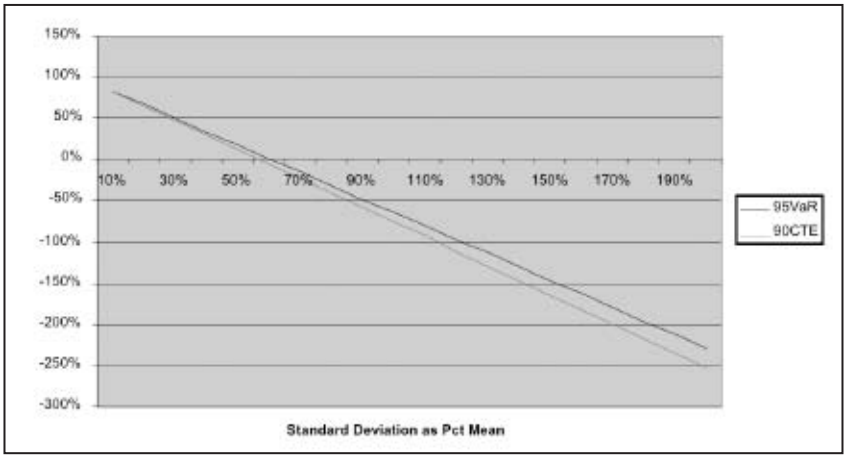

The conclusion here is that VaR and CTE are not necessarily significantly different if gains and losses are normally distributed. However, if gains and losses do not conform to a formulaic distribution, the difference is unpredictable. In those situations, CTE can provide significantly different information from VaR.

2. Tails and Number of Scenarios

At this point, we will switch from looking at a general loss distribution to a specific real model situation. The latest CTE discussions are revolving around variable products and guarantees in those products. As mentioned earlier, products and modeling results vary widely. However, significant insight into the workings of CTE can be observed from looking at the CTE of a simple portfolio of common stocks. The portfolio returns can be generated using several different scenario generation methods and the impact on CTE displayed.

The AAA committee report on C3 Phase II RBC suggests that a regime switching lognormal (RS2LN) model adequately captures the desired characteristics for the proposed CTE calculations. We are going to start with a lognormally distributed set of returns for a stock portfolio and eventually compare that to a RS2LN set of scenarios. For both sets, we will calibrate the models so that the mean is 11.9 percent and the standard deviation is 17.2 percent overall.

For the lognormal model, by stating those values, the parameterization is complete. There are only two parameters, mean and standard deviation. The ratio of standard deviation to mean is 1.45 percent and we would expect CTE90 for normally distributed losses to come in at 155 percent of the mean or an absolute value of 18.45 percent based on calculations behind the graph above. In fact, for our set of 1,000 scenarios, we get a CTE90 of 18.35 percent, which is 154 percent of the mean return.

For 20 sets of 1,000 lognormal scenarios, the CTE90 averaged 18.40 percent. In addition, we found that the CTE90 has range of 18.33 percent to 18.50 percent and that there is a 96 percent correlation between the CTE90 and the worst value of the 1,000 scenarios. That is a small cause for concern, since the worst value is probably the least dependable statistic of the entire set (Chart 4).

The ratio of the largest and smallest CTE values is less than 101 percent. That would suggest that the CTE90 calculation for 1,000 scenarios may be sufficiently reliable.

To look for a smaller number of scenarios that may be reliable, we took the same set of random values and divided it into 40 sets of 500 scenarios. The CTE90 results get much more diverse: (Chart 5)

The range of highest to lowest CTE value is 126 percent for these 40 sets of random returns. It would seem that the empirical evidence supports the idea that 1000 scenarios may be adequate, while 500 scenarios is not. The finding that there was high correlation between worst scenario and CTE90 value suggests that a user of CTE should examine the worst scenario to make sure that it is not causing an exaggeration of CTE value due to an extreme outlier in an insufficiently large scenario set.

3. Comparison to C1(CS)

The RBC requirements for life insurance companies already contain a component for common stocks. That amount is 22.5 percent for a fully diversified portfolio that qualifies for the lowest level of RBC based on the beta of the portfolio. If we calculate from the 20 sets of 1,000 random scenarios, the present value of the CTE90 result on year out, we get a result that ranges from 22.13 to 22.29. That looks to be very consistent with the 22.5 percent requirement that already exists.

That question and answer allow us to begin to understand the drastic difference between the risk of an investment holding like a common stock portfolio and an insurance product.

Chart 4—CTE90 for 1000 Scenarios

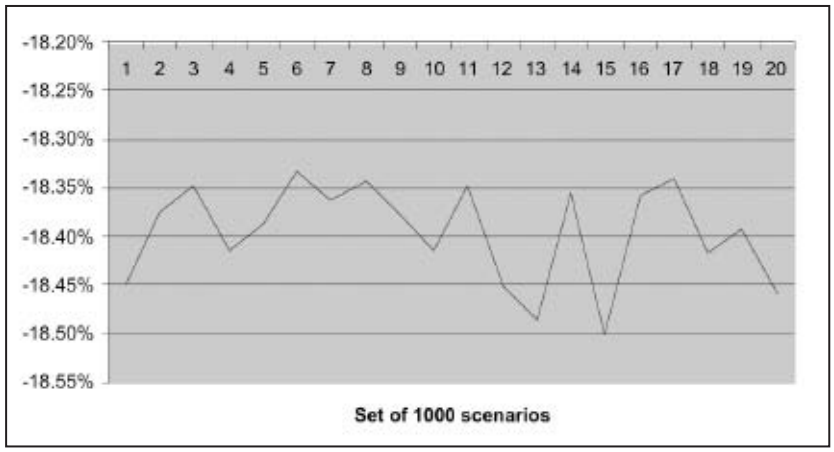

Chart 5—CTE90 for 500 Scenario Sets

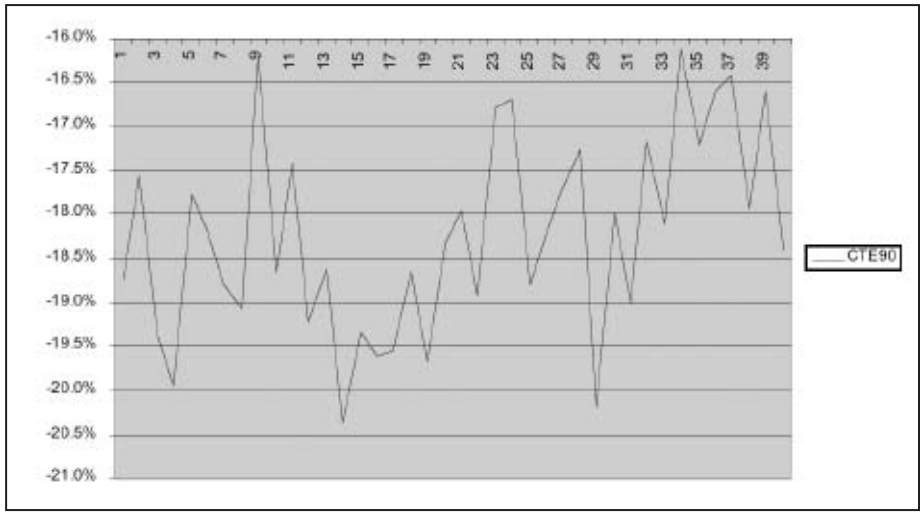

4. Modified CTE and Holding Period

When the CTE concept is applied to calculate C3 RBC, there are three differences from any of the above discussion. First, the calculation is projected over multiple years. For that purpose, we will use our 20 sets of 1,000 scenarios to now represent 1,000 scenarios of 20 years.

continued on page 36 ◗

Getting to Know CTE

◗ continued from page 35

Second, for RBC purposes, when calculating the CTE, a present value is calculated for each possible future period under each multi-year scenario and the worst present value is chosen. Finally, in calculating the CTE for any one year, scenarios, which would otherwise become a part of the average, are set to zero first if they are not in a loss position. In this section we will show the impact of each of those modifications to CTE related to our simple test case of a common stock portfolio.

First, let us look at the impact of extending the calculation to multiple years. The RBC for common stocks is equivalent to the CTE90 for a one-year holding period for our common stock model. One year is an appropriate standard for this comparison because there is no commitment on the part of an insurance company to hold stocks for any future period. Variable annuity products and the guaranteed benefits that are the concern of the C3 Phase II proposals do involve long-term commitments. For this example, we will use our stock model as if there were a long-term commitment to remain in the stock market. Most people have seen material during the 1990s bull market that emphasized the favorable impact of diversification over time. That material showed that risk decreases over time for a buy and hold stock investment strategy. However, the mechanics of the CTE measure produce the opposite effect. With CTE as the risk measure, risk appears to go up over time for several years, and then declines (Chart 6).

Why does that happen? While the cumulative standard deviation of returns gets smaller and smaller over time because there are ultimately up scenarios that moderate the down scenarios, the CTE gets worse because the make-up of the CTE90 is continually changing with each passing year. As the random scenarios generate new outliers each and every year, the scenarios that were moderated with the passage of time get pushed out of the CTE by these new outliers.

Chart 6—S&P 500 Normal Scenarios

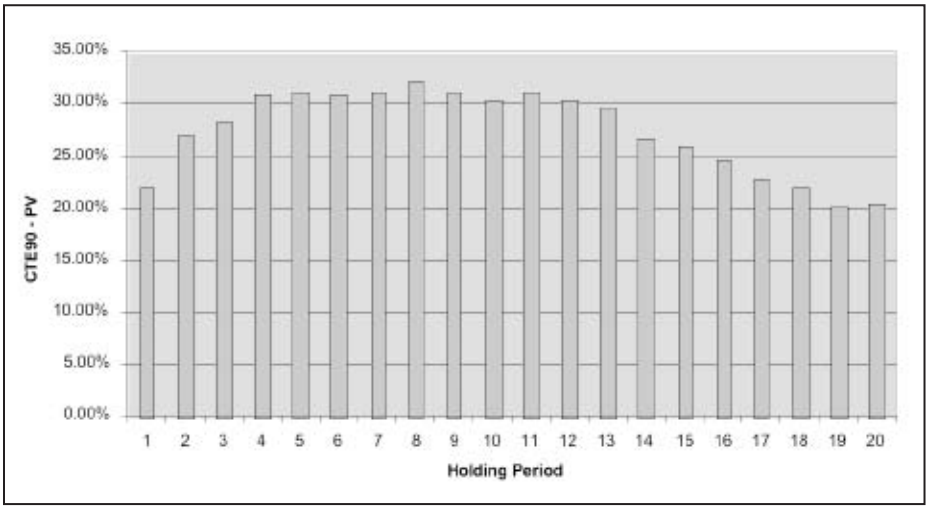

Table 3

Number of total years in worst 100	Worst 100 at 20 years	Next 100 Scenarios
0		32
1	5	10
2	4	14
3	3	6
4	6	1
5	2	7
6	5	3
7	4	7
8	5	4
9	6	4
10	5	2
11	9	5
12	3	2
13	5	0
14	9	0
15	5	2
16	5	1
17	3	0
18	6	0
19	7	0
20	3	0
	100	100

Table 3 shows that there are only three scenarios that were a part of the CTE90 for all 20 years and that five scenarios crept in only in the final year. In addition, the table shows that 68 of the next 100 scenarios ranked by final value of the stock fund, were a part of the CTE90 value for as many as 16 of the 20 years.

Second, we look at the impact of choosing the worst present value (Chart 7). The worst present value rule takes the forces that we mention above and maximizes their impact. Under this rule, the really extreme scenarios dominate no matter whether they are followed by better scenarios or not.

Finally, with the lowest PV adjustment, at least in this example, the need for the no-negative values adjustment is not needed. There are 344 scenarios to choose from in determining the one-year CTE and since no negative value is ever lost, that stable grows to over 600 possible negative values over the 20 years. That means that over 600 of the scenarios show that the stock portfolio is underwater at some time in the 20 years. That is not a statistic that the mutual fund and brokerage industry would have been excited to share when they were talking about how well stocks do when held for the long run (Chart 8).

In fact, if the no-negative adjustment is applied without the lowest PV adjustment, then there is an impact. With our scenario sets, there are less than 100 negative scenarios for holding periods of over 15 years.

This example shows that, as measured by CTE, there is negligible time diversification in this stock portfolio. With the modifications to CTE that are imposed on its use for RBC, extensions to holding period are measured to add significant risk.

continued on page 38 ◗

Chart 7—S&P 500 Normal Scenarios

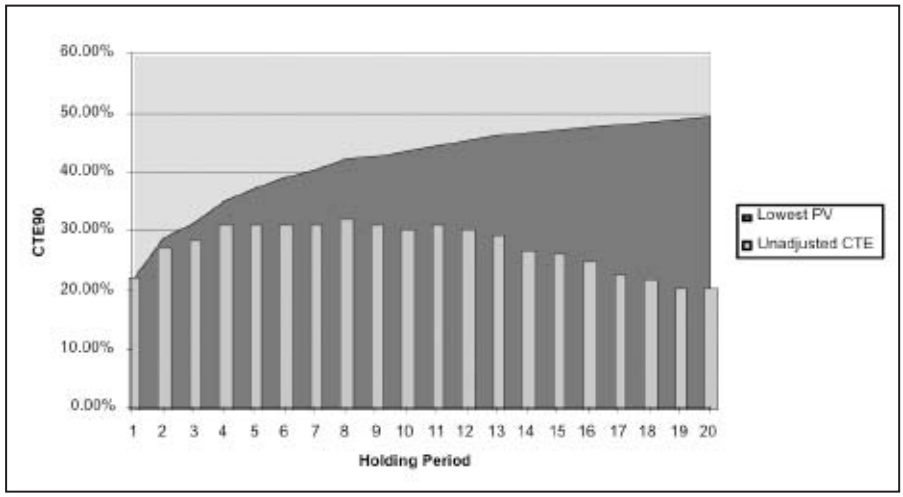

Chart 8—S&P 500 Normal Scenarios

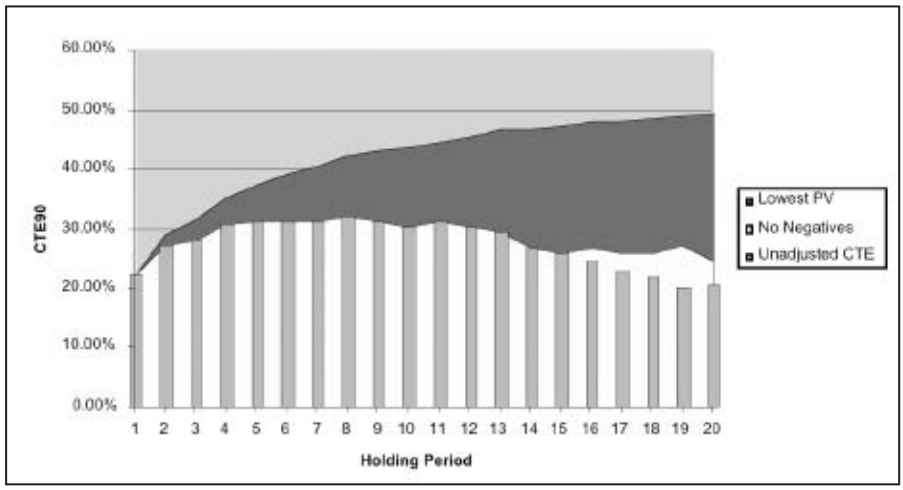

Chart 9—S&P 500 1000 Scenarios

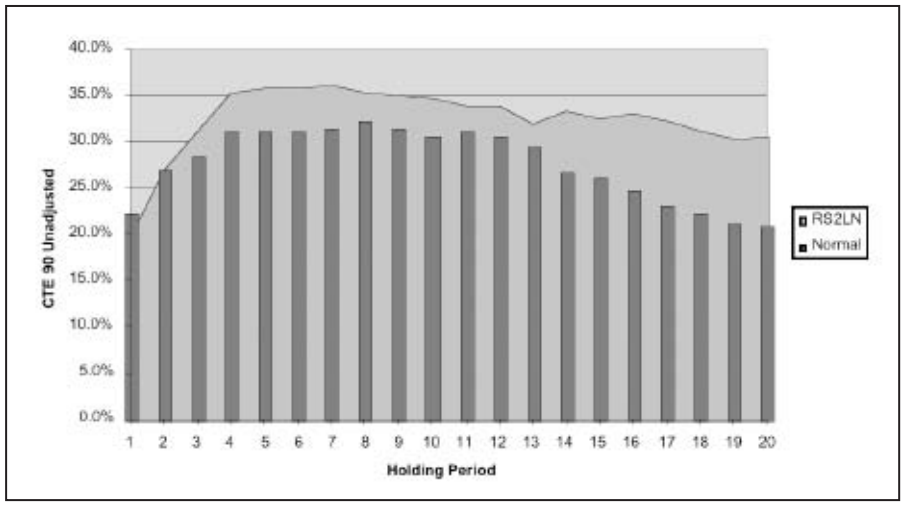

Getting to Know CTE

◗ continued from page 37

5. Regime Switching vs. Single Lognormal Scenarios and CTE

Now that we are good friends with CTE and its characteristics, we will look at the impact of different scenario generators. A regime-switching lognormal scenario generator (RS2LN) has been mentioned as a method that is sufficiently robust for the C3 Phase II calculations.

We created a set of RS2LN scenarios starting with the parameters in Mary Hardy's paper in the April 2001 *North American Actuarial Journal*. Over 20 years and 1,000 scenarios, the mean is 11.9 percent and the standard deviation is 17.5 percent. If we go back to our normal set of scenarios, we found that the one-year CTE on the return was -18.45 percent. In fact, our scenarios have a CTE90 for a one-year return of -19.98 percent. That means that the tail for these scenarios is somewhat more robust than the tail for normal scenarios.

Chart 10—S&P 500 1000 Scenarios

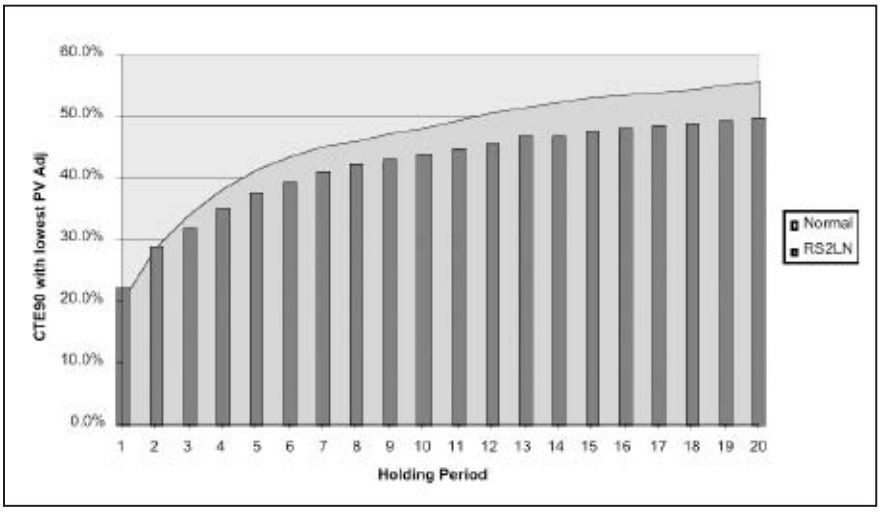

That more robust tail should and does translate into a somewhat larger CTE figure (Chart 9).

The difference is less severe after the adjustment for lowest PV. The 20-year holding period unadjusted RS2LN value is 45 percent higher than the normal value while the lowest PV value is only 12 percent higher (Chart 10). In fact, Chart 10 shows such a regular relationship between RS2LN and lognormal, that the extra effort does not seem to add any information. A slightly higher CTE (92 or 93) would give the same result under the lognormal scenarios as the CTE90 under RS2LN.

6. Conclusion

This examination of CTE has been performed butterfly fashion. I have skipped around looking at whatever caught my interest. If you are going to use CTE as a primary risk measure, then you should perform these tests and others to assure yourself that you understand why the CTE performs as it does. In this analysis, you should not get overwhelmed by the sheer volume of numbers that are involved. Even thousands and thousands of scenario results will eventually succumb to rigorous analysis.

Once mastered, CTE can be a very useful tool beyond its required use for RBC. As shown above, CTE is consistent with historical risk measures such as the C1 requirement for common stocks. CTE can also be used as a general risk measure and as criteria for replication (See "Why Write Variable Products When You Can Put the Money Directly into the Stock Market," Ingram & Silverman, *Risk & Rewards*, October 2003). ✦