Society of Actuaries

Article from:

Risks & Rewards

August 2014 – Issue 64

DOLLAR COST AVERAGING RISK

By Salil Mehta

Think of a security holding that one wants to sell over a period of time, being cognizant of the price risk associated with not selling the entire security immediately. For example, one can decide near a market top to sell securities over a half-year period, without properly hedging for the loss in security value during that time. Or one can imagine a hedging mechanism in reverse to, for example, acquire fixed amounts of a commodity across equal intervals of future time.

For this article, we use the baseline of an investment manager who wants to sell holdings of 30 shares of General Motors (GM) stock. The investment manager may consider selling the shares over time, for the purpose of reducing the liquidity risk that would come from selling all the shares at once on a somewhat arbitrary date. So we will try an approach of selling one share weekly, over 30 weeks. Now the traditional formula for understanding the cumulative risk for spreading out this sale over time assumes a *fixed* standard deviation (σ) for GM stock, for the entire 30 weeks. We will show that it also works by continuously summing the risk for the entire balance weekly, as this balance diminishes by one share weekly. And we'll discuss the drawbacks of this modeling approach. Later we'll explore how to think about a model, where σ instead increases or decreases by a fixed rate. For the baseline we start with a weekly standard deviation of 3 percent.

First we show the traditional formula for the extra risk (realized price variance) of trying to evenly liquidate a balance over a number of time periods n.

variance $= k*\sigma^2*q^2*(P_0)^2$

Where:

k	$= n/3*(1-1/n)(1-1/2n)$
q	= *number of shares*
P_0	= *initial price*

For the GM baseline example, we have:

k	$= 30/3*(29/30)(59/60)$
	$= 9.5$
$k*\sigma^2*q^2$	$= 9.5 * (3\%)^2 * 30^2$
	$= 7.7$
variance	$= 7.7*(P_0)^2$

This formula is also shown in popular risk books such as *Value at Risk* by Philippe Jorion.[1] In the formula one will notice that q, and P_0 are both constants, squared alongside the σ^2 that remains constant. The traditional formula shows that the focus for the total risk calculation is that it grows in proportion to the cube of n as n appears three times in the formula for k.

The reference for this formula is "declining-balance approach." But in this article we propose a newer theoretical methodology that helps a manager to gain a more intuitive feel for how total risk builds over lengthy trials (or in this case, lengthy amounts of time). We use instead an "incremental-time approach" that assesses the marginal contribution of each period to the total variance. Intuitively, risk is greatest for the last payment, not the first one. This insight can be applied to price a variety of term-risk contracts (e.g., how much capital to reserve away to hedge the risk of an expected payment such as an inheritance or bonus, or a large expense such as estimated taxes or college tuition, or the risk of systematic liquidation of a guaranteed investment contract).

CONTINUED ON PAGE 18

Figure 1

Theoretical marginal risk contribution to fixed total risk

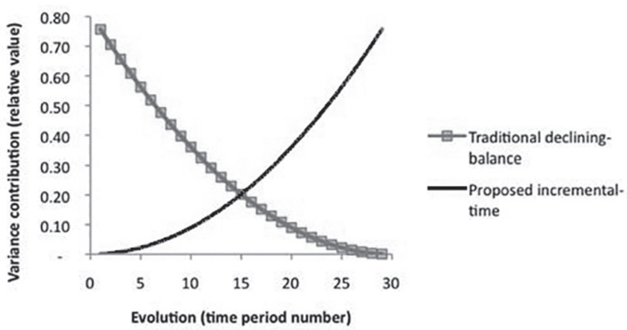

We'll be able to visualize through the illustration below that the logic of the traditional formula is that it assumes individual weeks are all contributing equally to the total risk. For this illustration we show the newer proposed approach as well, and we again use the baseline practical weekly σ assumption 3 percent (roughly 20 percent annualized).

The proposed mathematical method has the advantage of forcing one to appreciate the incremental risk, associated with a marginal increase in the time ***n***, of a balance liquidation. On the other hand, the traditional approach seems to imply risk is lowest at the end of the liquidation period, at the latest time it discretely approximates as ***n***-1, when in fact it is of highest risk then.

Let's explore what makes these mathematical properties work. It comes down to the formula for the variance of a series of independent and identically distributed, normal random variables. This is different from the approach Newton Bowers uses in *Actuarial Mathematics* to price annuities and life insurance. In the article here, the payment amounts are variable but the time period is fixed.

Let's explore the manipulation of the variance mathematics further. Suppose the variance per period of a single share is σ^2, and we start with ***n*** shares to liquidate. By the ordinary properties of variance:

Total variance =
$[(n-1)/n]^2 \sigma^2 + [(n-2)/n]^2 \sigma^2 + \ldots + [1/n]^2 \sigma^2$

The theoretical variance associated with the total risk is relative to the fixed sizes of ***n***-1 shares, to one share. Or σ^2/n^2 times the sum of: $(n-1)^2+\ldots+3^2+2^2+1^2$. We make one adjustment partway into the solution, since we assume no marginal volatility contribution associated with the first immediate share sale, from the total relative size of ***n*** shares. We also now algebraically rearrange this expression and demonstrate the flexibility of its usage.

Start with the special geometric growth series:
$1^2+2^2+3^2+\ldots+n^2$
$= n(n+1)(2n+1)/6$
$= (2n^3+3n^2+n)/6$

We can substitute (n-1) for n, and the sum:
$1^2+2^2+3^2+\ldots+(n-1)^2$
$= (2n^3-3n^2+n)/6$
$= n(n-1)(2n-1)/6$

And the constant proportional weights imply 1/n² times each term above:
$n(n-1)(2n-1)/6n^2$
$= n/6*[(n-1)/n][(2n-1)/n]$
$= n/3*[(n-1)/n][(2n-1)/2n]$
$= n/3*(1-1/n)(1-1/2n)$

Given the linear connection between the sum, and the sum of these variances, we can reconstruct and describe this final formula using the proposed approach. See the traditional declining-balance approach on the left of the illustration in Figure 2 at the top of page 19. Then see the proposed incremental-time approach, which comes to the same total amount, as shown on the right side of Figure 2.

INTUITIVELY, RISK IS GREATEST FOR THE LAST PAYMENT, NOT THE FIRST ONE.

To be sure, the first and last vertical bars on the left of the illustration (top, right), for the traditional approach, the variance contribution is:
$((n\text{-}1)/n)^2*\sigma^2*q^2*(P_0)^2$, and $(1/n)^2*\sigma^2*q^2*(P_0)^2$, respectively.

While the first and last vertical bars on the right of the illustration (top, right), for the new proposed approach, the variance contributions are reversed:
$(1/n)^2*\sigma^2*q^2*(P_0)^2$, and $((n\text{-}1)/n)^2*\sigma^2*q^2*(P_0)^2$, respectively.

And both triangular bars again sum to $9.5*\sigma^2*q^2*(P_0)^2$, or the $7.7(P_0)^2$ we showed in the initial math for the GM baseline example.

Now the second advantage of the proposed incremental-time approach, besides the intuition of incremental variance per extra unit of time n, is that one can also disaggregate and select the individual terms for hedging. For our 30-week GM stock example, we can propose that the security doesn't maintain a fixed σ, but we can instead assume an example later where the σ may increase by 4 percent weekly instead of 0 percent. And we may want to understand the value to offset the risk of specific terms to manage liquidity (e.g., to offset a tax payment, or hedge a special dividend announcement).

Now empirical evidence shows that markets are not always a fair random walk. Sometimes there is unusually strong serial correlation, similar to that which we have seen over 2013 and year-to-date in the U.S. stock market. But keep in mind that this autocorrelation would bias the results only slightly for the traditional baseline risk formula as well. And the underestimation of risk by not considering it makes understanding of the newer proposed risk approach that much more valuable.

Here too the traditional declining-balance approach could not handle these additional illustration requirements, even though the mathematics seemed fairly benign at the start of this note. This can be shown more completely through a stochastic *simulation* model. See the illustration above (bottom, right) where we simply simulate the baseline example, where we see the effect of thinking about the amount of risk relative to when the payment is made. Then in the illustration in Figure 4 on page 20, we allow σ to vary over time. Notice the dashed lines (since there we show simulations) are exponentially growing, but are far more stable. This shows the disadvantage of using the traditional declining-balance approach versus the cleaner proposed approach.

FIGURE 2

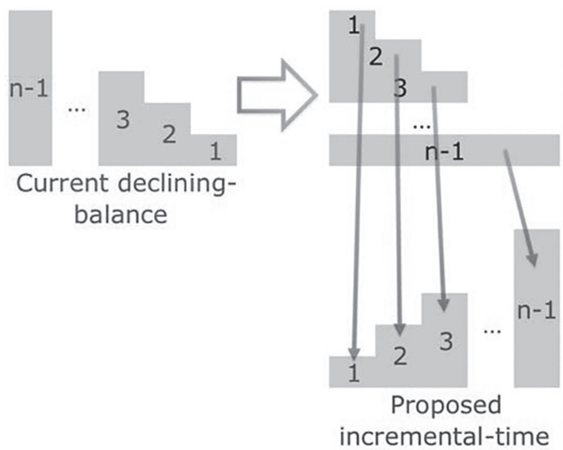

Figure 3 Using Fixed σ
Simulated liquidity risk per future payment, using both approches

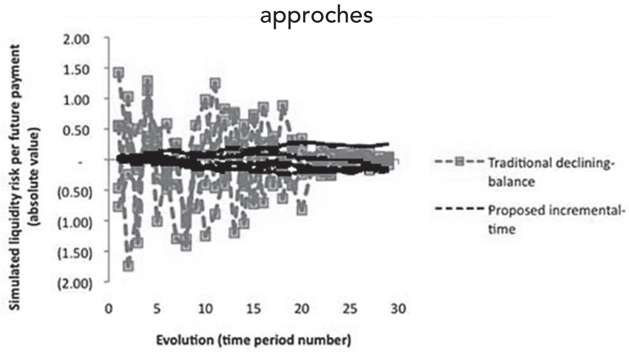

CONTINUED ON **PAGE 20**

As noted above, if the σ doesn't *vary* at all, then these collapse to our normal baseline, and both of the simulation approaches (and both theoretical approaches) all agree. Total variance for the traditional declining-balance approach, or the proposed incremental-time approach, both also equal 7.7, which we would multiply of course by $(P_0)^2$. The simulation runs many samples, and the sums of the first five shown in the chart are 0.4, -2.9, 5.1, -3.4 and -3.5. The variance among this broader sample is 7.7 regardless. But, for example, in the new proposed approach we can see a steadier and more intuitive build-out evolving over the liquidation time, for the five sample values.

But as we change σ, the additional variance from the proposed approach begins to grow in a convex fashion. At 2 percent (with n=30), there is more than a 50 percent difference in total risk estimate. And at 4 percent, as we show in the chart below, there is about a 175 percent difference. Notice first that the axis scale is now enlarged versus the

Figure 4 Using Increasing σ

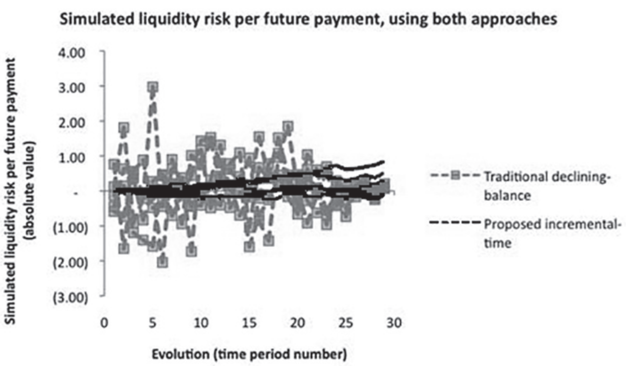

Note change of scale versus Figure 3

simulated baseline chart. Of course again this new simulation runs many samples, and the sums of the first five shown in the chart are -2.1, -0.1, 4.7, 6.8 and 0.7. The variance among this broader sample is 16, or greater than 100 percent difference from the 7.7 (baseline approach). For the proposed approach the corresponding totals are -3.8, 0.1, 7.2, 10.3 and -0.6. The variance among this broader sample is 36, or greater than 100 percent difference from the 16 (the traditional approach).

Bear in mind that an aggressive 4 percent weekly increase in the σ was illustrated above. In most practical cases we would see assumptions about one-half of this, and the differences would be about one-third of what is illustrated above. We show this broader range to illustrate the differences in variance approximation that exist between these two approaches, which are designed to answer different types of risk management questions against changing σ regimes.

To summarize, the proposed mathematical approach of building total variance from the incremental contribution of each payment comes with no downside versus the traditional approach. But the proposed approach offers a cleaner and more reliable insight into time-specific risk contribution, and allows one to consider the real-world usefulness of varying the σ risk over the uniform liquidation period.

THE PROPOSED APPROACH OFFERS A CLEANER AND MORE RELIABLE INSIGHT INTO TIME-SPECIFIC RISK CONTRIBUTION, AND ALLOWS ONE TO CONSIDER THE REAL-WORLD USEFULNESS OF VARYING THE RISK OVER A UNIFORM LIQUIDATION PERIOD.

For further details, visit *https://sites.google.com/site/statisticalideas/home/term-risks-math* for "Term Risks Math" on the free "Statistical Ideas" resource portal for academics and practitioners alike. It lets the users explore the traditional and proposed risk methods, using the fixed or varying σ assumptions, and in theoretical form or simulation. Instructions are provided on the Web portal.

ENDNOTES

[1] Page 344, third edition, McGraw Hill.

Salil Mehta is a statistician and risk strategist, as well a lecturer at Rutgers and Georgetown University. He served for two years as the group Director of Analytics, in the U.S. Department of the Treasury for the Administration's $700 billion TARP program. Salil is also the former Director of the Policy, Research, and Analysis Department in the Pension Benefit Guaranty Corporation. He is a CFA and candidate member of the SOA. He is the author of Statistics Topics (http://statisticalideas.blogspot.com/p/book.html) which provides an interesting discussion of various statistical applications.

2014 SOA Annual Meeting, Orlando Fla., Oct 26-29

INVESTMENT SECTION-SPONSORED SESSIONS

38 L Economic Capital: Key Modelling Considerations
77 Investment Section Hot Breakfast
87 PD Smart -- and Dumb -- Things About Smart Beta
101 PD Use of Efficient Frontiers in Strategic Asset Allocation Analysis
118 PD Variable Annuity and Indexed Annuity Risk Management Strategies
133 PD Tail Risk Hedging
172 PD Investment Strategies and Alternative Investments in Insurance and Pension Portfolios
183 PD Predictive Modeling for Actuaries: Machine Learning, Predictive Modeling and Insurance Risk Management

Looking forward to seeing you at the Section breakfast!
And don't forget the 2015 Investment Symposium early next year.