Why is the FRTB Expected Shortfall Calculation Designed as It Is?

Greg Hopper | May 23, 2023

The Fundamental Review of the Trading Book (FRTB) will be the new market risk capital standard for banks. It is slated for implementation in the E.U. and U.K. by Jan. 1, 2025 and expected to be applied in the United States by that same date. Having taken years for the Basel Committee to develop, the FRTB is designed to address problems with the existing Value at Risk (VaR)-based market risk capital measures under the international Basel capital standard (often referred to as "Basel II.5").

The global financial crisis revealed fundamental problems with the estimation of VaR that led to an underestimation of market risk and therefore appropriate market risk capital.[1] The FRTB introduces a new measure—Expected Shortfall (ES)—to replace VaR. The FRTB also establishes very stringent eligibility requirements for using ES to calculate the market risk associated with respect to a particular asset. Assets that do not meet these requirements must either undergo a stress test or be subject to an extremely conservative standardized approach for determining market risk. Thus, the FRTB employs non-simulation approaches to address problems with VaR that cannot be resolved by ES.

The implementation of the new standard will dramatically increase market risk capital requirements. The most recent Basel III Monitoring Report[2] estimated that the weighted-average market risk capital requirement resulting from replacing the VaR capital standard with FRTB is expected to increase by 63.2 percent for Group 1 Banks[3] and 69.2 percent for the GSIB subset of Group 1, with the increase largely felt in the U.S.

Meanwhile, under the auspices of Dodd-Frank, the Federal Reserve developed a market risk stress test, the Global Market Shock (GMS), that also addresses the underlying problems in the VaR calculation. Since 2020, the GMS has been a component of the Stress Capital Buffer, which is part of each large bank's minimum capital requirement. The Stress Capital Buffer calculates how much a bank would lose under extreme stress, and requires the bank to hold capital equal to that loss, in addition to a 4.5 percent capital charge, any GSIB surcharge and any countercyclical charge.

For banks subject to the GMS, the total market risk capital is calculated by adding the GMS-generated market risk capital charge to the market risk capital charge under the Basel II.5 VaR standard. When FRTB is implemented in the U.S., total market risk capital will be calculated by adding the market risk capital charge from the GMS to the market risk capital charge under FRTB.

This change prompts the question of whether the GMS and FRTB should be added together in the same way that the GMS and Basel II.5 were added together, which means determining whether they are capitalizing different risks (as with GMS and Basel II.5) or some of the same risks. The GMS shocks (and the risks they capture) are simple to understand, but the FRTB methodology is filled with arcane formulas that make it hard to see the

[1] See for example "An Evaluation of Bank VaR Measures for Market Risk During and Before the Financial Crisis," Finance and Economics Discussion Series, available at https://www.federalreserve.gov/econres/feds/an-evaluation-of-bank-var-measures-for-market-risk-during-and-before-the-financial-crisis.htm.

[2] See "Basel III Monitoring Report," September 2022.

[3] Group 1 banks are those that have Tier 1 capital of more than €3 billion and are internationally active.

motivation behind the calculation or what risks it is designed to capture. For policymakers who are considering updates to the GMS, this post is designed to help de-mystify the FRTB Expected Shortfall calculation by working through a simple example step-by-step, noting at each point what risks are being captured. It is worth noting, however, that the FRTB introduces new problems with market risk capital estimation due to its use of non-simulated approaches for cases that cannot be covered by ES. We also identify these new problems in this post.

In a follow-up post, we will continue the analysis in this post to show that the FRTB in large part duplicates the measurement of the risks the GMS is designed to cover, thereby creating additional market risk capital requirements for trading that are not commensurate with the underlying risk. That would have major implications for the liquidity and robustness of U.S. capital markets, which we have previously documented in "The Global Market Shock and Bond Market Liquidity."[4] Although the FRTB is a significant new addition to the market risk capital framework, the GMS has not been meaningfully updated since its inception. In the follow-up post, we will make some suggestions on how the GMS can be updated to work better with the FRTB.

How is VaR Calculated?

To illustrate a VaR calculation, we will assume a simple portfolio: a long position of $100 million in the Dow Jones Small Cap Index, a $500 million long position in a five-year credit default swap (CDS) on the investment-grade CDS index, and a $500 million long position in a five-year credit default swap on the high-yield CDS index. For both CDS positions, the bank receives the premium and pays out on a default. This hypothetical portfolio will rise in value if small-cap stocks go up, and also rise in value if the chances of default for investment-grade or junk bonds increase.

Suppose we want to estimate the 99th percentile one-day VaR, i.e., the loss such that 99 percent of the time it will be lower than the calculated VaR and 1 percent of the time it will be equal to or greater than the calculated VaR. There are many methods that can be used to calculate VaR, but for purposes of illustration we will use a common, simple method—the "historical" method.

To employ the historical method, we first need to specify the historical data to use for the estimation. Banks use between one and four years of past data for this calculation; for simplicity, we will use the last 500 days of daily data, specifically the period between Feb. 24, 2021 to Feb. 16, 2023. On each day, we calculate the daily change in value of the three assets and sum them to get the portfolio change. We then sort the portfolio results from largest to smallest losses. The sorted losses allow us to read off our VaR estimate for the 99th and other percentiles. Exhibit 1 shows the results, with VaRs at the 99 percent, 98 percent and 97 percent shaded in yellow. Note that the portfolio of losses for the 500 days has been sorted starting from the worst loss, which occurred on June 13, 2022. The 99 percent VaR, which occurred on May 5, 2022, is $8.8 million, since 1 percent of the losses, i.e., 1 percent of the 500 daily portfolio changes, or five losses, are equal to or greater than $8.8 million. Similarly, the 98 percent VaR is $7.1 million, since there are 10 losses equal to or greater than $7.1 million.

Basel capital is calculated using a 10-day VaR

In Exhibit 1, we calculated a one-day VaR. However, Basel capital is based on a 10-day VaR, which is designed to account for the potential inability of banks to hedge or trade for as long as ten days because of market illiquidity. Do we then need to estimate VaR using 10-day returns? In practice, banks have avoided the need to use 10-day returns by multiplying one-day VaR by $\sqrt{10}$, which is permitted by the Basel capital rules. This scaling method is valid as long as daily asset returns follow the same normal distribution and no day's return influences any other day's return[5].

[4] This note is available at https://bpi.com/the-global-market-shock-and-bond-market-liquidity/.

[5] In practice, both of these assumptions are not true. Returns are not normally distributed in general and returns today depend on returns over past days. However, the scaling method can often be a useful approximation to scale up VaR over different time periods.

Thus, under Basel capital rules, we do not need to re-estimate VaR using 10-day returns but can rather calculate 10-day VaR as $\sqrt{10}\,(8.8) = \$28\ million.$

Exhibit 1

Date	Equity Return	IG CDS Return	HY CDS Return	Equity Loss	IG CDS Loss	HY CDS Loss	Portfolio Loss	VaR Percentile
6/13/2022	-4.9%	7.6%	8.1%	-4,887,644	-1,241,941	-4,490,356	-10,619,942	
6/16/2022	-4.7%	8.3%	7.7%	-4,726,354	-1,363,591	-4,288,648	-10,378,599	
9/13/2022	-3.8%	8.5%	9.2%	-3,777,086	-1,396,700	-5,118,910	-10,292,776	
5/18/2022	-3.7%	7.9%	6.9%	-3,741,655	-1,292,044	-3,825,910	-8,859,609	
5/5/2022	-3.8%	6.4%	7.1%	-3,775,348	-1,052,569	-3,947,279	-8,775,196	99%
2/25/2021	-3.2%	8.5%	7.2%	-3,220,623	-1,389,995	-3,996,182	-8,606,740	
8/26/2022	-3.1%	6.3%	6.2%	-3,091,003	-1,029,247	-3,453,861	-7.574.111	
11/26/2021	-2.7%	8.6%	6.2%	-2,725,290	-1,410,772	-3,424,114	-7,560,177	
12/3/2021	-1.5%	10.7%	7.6%	-1,520,030	-1,748,451	-4,196,682	-7,465,174	
4/26/2022	-3.0%	6.6%	5.5%	-2,975,942	-1,082,534	-3,077,045	-7,135,521	98%
6/10/2022	-3.0%	4.8%	6.1%	-2,962,306	-783,264	-3,361,452	-7,107,022	
3/7/2022	-3.6%	4.4%	5.0%	-3,604,535	-718,242	-2,777,958	-7,100,734	
5/9/2022	-4.3%	2.3%	3.4%	-4,301,221	-375,765	-1,872,026	-6,549,012	
12/15/2022	-2.3%	6.0%	5.5%	-2,297,690	-990,085	-3,091,905	-6,319,700	
9/23/2022	-2.2%	3.5%	6.1%	-2,234,223	-570,272	-3,384,529	-6,189,024	97%

VaR calculation problems and FRTB solutions

Below are problems with VaR (and thus with Basel II.5) that the Fed sought to correct by adding a capital surcharge through the GMS and that the Basel Committee sought to correct by replacing Basel II.5 with the FRTB.

Problem 1: VaR does not capture rare large losses effectively

VaR can easily fail to capture rare, large losses. For example, suppose we engage in an option-like trade that pays a small gain most of the time, but very rarely results in extremely large losses. For example, over the past 500 days, it might turn out that on 497 days the asset gained $10,000, but on three days it lost $100 million dollars. The 99 percent VaR would be determined by the fifth worst loss, which is a gain of $10,000. The $100 million loss would be missed.

Problem 1: FRTB solution

To correct the potential underestimation in VaR, the Basel Committee decided to replace VaR with an Expected Shortfall (ES) measure. The idea behind ES is to average the losses above VaR. For example, if we calculated one-day ES at the 99 percent level, we would average the four losses above the fifth worst loss, which is the 99% VaR for 500 days of data. In our hypothetical example, the one-day ES would be the average of $10,000 and three $100 million losses, or approximately $75 million.

Thus, in sum, while VaR (and thus Basel II.5) tells us how bad losses will be for a given probability, ES tells us the average loss if things do in fact get that bad – that is, if losses are in the tail beyond the historically likely realm of loss. ES is a better measure than VaR when the risk of loss has "fat tails," since ES averages the risk in a fat tail.[6] In this sense, it could not be more different from VaR.

Rather than define ES at the 99 percent level, the Basel Committee specified the ES at the 97.5 percent level because a 97.5 percent ES will be very similar to a 99 percent VaR so long as the losses are approximately normally distributed.[7] Even when losses are not normally distributed, ES at 97.5 percent can still be a reasonable approximation to the 99 percent VaR. For example, if we calculate one-day ES at the 97.5 percent level from Exhibit 1, we would average the losses between $7.1 and $10.6 million, yielding a one-day ES of $8.5 million, close to the 1-day 99 percent VaR estimate of $8.8 million.

Problem 2: The future will not be like the past

In exhibit 1, we estimated VaR based on what values it took over the past 500 days. The implicit assumption is that so long as the future is essentially like the past, the historical VaR will be a good estimate of VaR going forward. However, it is rarely the case that the future is like the past in financial markets. In especially turbulent periods, such as during 2008-09, VaR estimated on past data can significantly underestimate the frequency of future large daily losses.

One way to resolve this problem is to estimate an historical VaR over a particularly volatile historical period-known as a "stressed VaR" approach. So, instead of looking at the most recent historical period, we look at the worst-ever historical period. Basel II.5, in place currently, requires that a stressed historical VaR be estimated over a very volatile one-year period. If we choose Aug. 1, 2008 to July 29, 2009 as a stressed 250-day historical period and repeat the calculation, we would obtain a one-day stressed VaR estimate of $16.2 million. Scaling to 10 days, the stressed VaR estimate would be $\sqrt{10}\,(16.2) = \$51.2\ million$, about twice the *un*stressed VaR estimate. Under Basel II.5, capital is estimated by adding three times the 10-day VaR to three times the stressed 10-day VaR calculation.

Problem 2: FRTB solution

The FRTB adopts the same solution as Basel II.5 by requiring that ES be calculated during a stressed period. Since ES is designed to be a stressed measure, the FRTB does not permit the scaling method to be used to calculate 10-day ES but rather requires that ES be calculated directly using 10-day returns. To illustrate how ES would be calculated under the FRTB requirements for 10-day returns using our sample portfolio, we will employ the 250-day stressed market period between Aug. 1, 2008 and July 29, 2009. Exhibit 2 shows that we can estimate Portfolio ES at the 97.5 percent level using 250 days of data[8] by averaging the top five 10-day losses, shaded in yellow, yielding an ES of $51 million. Note that for this portfolio and dataset, the 97.5 percent 10-day ES of $51 million is the same as the 99 percent scaled 10-day stressed VaR estimate of $51 million. The scaling method is a very good approximation to using actual 10-day returns in this case.

[6] Losses can have "fat tails," implying that rare large losses are more common than implied by the normal distribution. When losses have fat tails, ES can be much larger than VaR.

[7] The normal distribution is a standard distribution that arises theoretically in probability and statistics but the evidence that financial variables are not normally distributed is very strong.

[8] The 10-day return periods, taken over a 1-year horizon, necessarily require that the returns overlap, implying that they are not independent. Using overlapping returns does not bias the results, but it effectively reduces the nominal 1-year sample size significantly.

Exhibit 2

Date	Equity Return	IG CDS Return	HY CDS Return	Equity Loss	IG CDS Loss	HY CDS Loss	Portfolio Loss
10/10/2008	-26.0%	34.8%	44.7%	-25,994,113	-5,657,175	-23,842,218	-55,493,505
10/9/2008	-28.2%	24.8%	40.0%	-28,224,366	-4,052,335	-21,420,314	-53,697,015
10/15/2008	-25.2%	14.6%	41.1%	-25,174,251	-2,389,515	-21,981,915	-49,545,681
11/20/2008	-22.2%	43.4%	35.3%	-22,167,689	-7,036,041	-19,003,334	-48,207,064
11/21/2008	-19.7%	41.5%	37.7%	-19,706,067	-6,742,417	-20,232,911	-46,681,395

Average of losses in yellow = Expected Shortfall = $51 million

Problem 3: All asset classes do not have the same 10-day liquidation period

Because the Basel II.5 VaR is calculated at a 10-day horizon, the VaR scaled to 10 days implicitly assumes that all assets can be liquidated over a 10-day period. The 10-day ES calculation of $51 million also assumes all assets have a liquidation period of 10 days. However, widely varying liquidity was observed for different asset classes during the financial crisis.

Problem 3: FRTB solution

The FRTB uses different liquidity assumptions for different asset classes, with a floor of 10 days. To illustrate the solution, Exhibit 3 below lists some of the different liquidity horizons for different asset classes specified in the FRTB. ES for assets that depend on large cap equities can be computed using the 10-day floor, but ES for assets that depend on small cap equities must calculated using a 20-day horizon. However, IG credit spreads require a 40-day horizon, four times the floor of 10 days.

Exhibit 3

Risk Factor	Days	Risk Factor	Days
IG Credit Spreads	40	Other Equity	60
HY Credit Spreads	60	EUR, GBP, USD Rates	10
Large Cap Equity	10	Standard FX Pairs	10
Small Cap Equity	20	Energy	20

The requirement to calculate ES at horizons greater than 10 days for different asset classes immediately presents a practical problem. With only 250 days of stressed data, it is not feasible to use 20-day, 40-day or 60-day returns, as the underlying data set is too small. Although the FRTB does not allow the scaling method to be used to calculate 10-day ES, it does allow the scaling method to be employed to estimate ES over periods longer than 10 days.

Implementing the scaling method under FRTB is more complicated since the underlying assets classes in the 10-day ES must be scaled to different horizons, rather than to a common 10-day horizon. For example, in our example portfolio the equity exposure must be scaled to 20 days (since it is a small cap index), the IG exposure must be scaled to 40 days and the HY exposure to 60 days.

The FRTB methodology does all scaling on 10-day ES estimates. In order to do the scaling, it will be necessary to estimate 10-day ES for each combination of risk factors that share a liquidity horizon. Since all three risk factors share a risk horizon over the first 20 days, 10-day ES using all three risk factors must be estimated and then scaled to 20 days. We already performed that 10-day estimate and found that it was $51 million. IG and HY share a risk horizon between 20 and 40 days, so 10-day ES that includes both IG and HY (but not equities) must be estimated and then scaled for an additional 20 days. Between 40 and 60 days, only HY is covered by the risk horizon and thus we must also estimate 10-day ES for HY only and scale it an additional 20 days. Then we must combine all these re-scaled estimates.

To illustrate how the scaling process works, we calculate 10-day ES for the combinations of the asset classes for our sample portfolio, shown in Exhibit 4. We also include ES calculated for equity only for reasons that will be explained later.

Exhibit 4

Equity, IG, and HY included	($mil, 10-day)	IG and HY included but equity excluded	($mil, 10-day)	Only HY included	($mil, 10-day)	Only EQ Included	($mil, 10-day)
$ES_{EQ,IG,HY}$	51	$ES_{IG,HY}$	26	ES_{HY}	22	ES_{EQ}	25

Exhibit 5 demonstrates the full ES calculation in which scaling is applied beyond ten days. As can be seen, the scaling method is the same as used for Basel II.5, except that it is done over 10-day rather than one-day periods. So, for example if we want to scale the 10-day ES estimate of 51 in Exhibit 4 to 20 days, we multiply 51 by the square root of two, i.e., $\sqrt{2}(51)$, just as in Basel II.5. However, unlike Basel II.5, we are adding different values of ES at different liquidity horizons. The ES provides a more general scaling method for this case: we square each term (e.g., $\left[\sqrt{2}(51)\right]^2 = 2(51)^2$), add them together, and then take the square root to get the final result, as illustrated in Exhibit 5.[9]

This calculation includes diversification of risk across equities and CDS, so we term the result the diversified ES.

[9] The technical reason for the ES scaling method is the same as used for Basel II.5 scaling. If returns are approximated to be from a normal distribution with a mean of zero, then $ES = 2.3378\sigma$, where σ is the standard deviation of the 10-day return. If returns are also independent across 10-day periods, then we can use the well-known formula for the variance of a sum of random variables to calculate the variance across multiple 10-day periods by adding variances of 10-day periods together. So, for example, if we have two 10-day variances σ_A^2 and three 10-day variances σ_B^2, then the 50-day variance would be $\sigma_{50}^2 = 2\sigma_A^2 + 3\sigma_B^2$. Then we take the square root of σ_{50}^2 and multiply by 2.3378 to get ES. Exhibit 5 illustrates this calculation.

Exhibit 5

Days 0-20	Days 20-40	Days 40-60	ES^2_{total}	$ES_{diversified}$
$2(51)^2$	$+\ 2(26)^2$	$+\ 2(22)^2$	= 7,522	$87 million
We have two 10-day periods in which equities, IG, and HY share a common liquidity horizon, so we scale the 10-day $ES^2_{EQ,IG,HY}$ by two	We have an additional two 10-day periods in which IG and HY share a common liquidity horizon, so we scale the 10-day $E^2_{IG,HY}$ by two. Equities drops out since its liquidity horizon terminated after 20 days	We have an additional two 10-day periods in which only HY is covered by the liquidity horizon, so we scale the 10-day ES^2_{HY} by two. IG drops out since its liquidity horizon terminated after 40 days	Sum the numbers in the three-time buckets	Take the square root of 7,522 to get the final ES estimate

Problem 4: We cannot always rely on diversification to reduce risk

In normal market conditions, risk factors do not generally move in the same direction. In a diversified portfolio, one asset might experience a large loss, but the other assets might just see small losses, or even offsetting gains. The Basel VaR calculation as well as the ES calculation detailed above allow diversification benefits across asset classes that lower both VaR and ES. However, during very stressful market conditions, diversification can break down, with large losses in one asset class not necessarily mitigated by smaller losses or gains in another asset class.

Problem 4: FRTB solution

To account for the potential breakdown in diversification benefits, the FRTB requires additional calculations of ES by asset class only, to eliminate the benefits of diversification across asset classes. To exclude diversification across the equities and CDS asset classes in our example portfolio, we would calculate ES for equities and CDS separately and then add them together. Exhibits 6 shows the calculations.

Exhibit 6

Days 0-20	Equities	Days 0-40	Days 40-60	CDS	$ES_{undiversified}$
$2(25)^2$	$35 million	$4(26)^2$	$+\ 2(22)^2$	$61 million	= $96 million
We have two 10-day periods for equities, so we scale the 10-day ES^2_{EQ} by two	Take the square root to get full ES calculation for equities for 20 days	We have four 10-day periods in which IG and HY share a common liquidity horizon, so we scale the 10-day $E^2_{IG,HY}$ by four.	We have an additional two 10-day periods in which only HY is covered by the liquidity horizon, so we scale the 10-day ES^2_{HY} by two.	Sum the ES for CDS over 60 days and take the square root	Add $35 million for equities and $61 million for CDS. Note that we did not give any benefit to correlation between equities and CDS

Finally, the FRTB rules specify that the ES used in the capital formula is the average of the diversified ES and the undiversified ES, $87 million from Exhibit 5 and $96 million from Exhibit 6, respectively:

$$ES_{Total} = \frac{87 + 96}{2} = \$91.5\ million$$

Problem 5: VaR cannot be estimated reasonably when the risk factors are relatively illiquid

In our example VaR and ES calculations, we had over 15 years of daily data and so could calculate VaR under current or past market conditions, including during the market stress period of 2008-09 as well as during the stressed conditions in 2020. In practice, there is not adequate daily data for every risk factor, implying that VaR may be misleadingly estimated.

Problem 5: FRTB solution

The FRTB solves this problem by requiring that all risk factors that are included in the ES pass specified liquidity and other tests to be modelable in ES. If a risk factor does not pass the test, it must be removed from the ES calculation and subject to a stress test with a liquidity horizon at least as long as defined for the ES calculation itself. The result of that stress test is added to the capital formula separately. Other non-modelable assets that do not qualify for the stress test must be estimated using a standardized approach.

New problems created by the FRTB

Although the FRTB largely resolves the five problems with VaR discussed above, it creates some new problems that should be resolved before it is implemented. These problems arise because the FRTB combines modeled and non-modeled approaches for estimating market risk capital and imposes stringent tests to use the model-based ES approach. These new problems produce an unjustified increase in market risk capital for US firms caused by overlap with the GMS, unnecessary volatility of market risk capital, and incentives to lower portfolio diversification. The first problem can be resolved by updating the GMS, and the other two have very simple solutions.

Overlap with the GMS

As previously discussed, risk factors that do not meet the requirements for treatment in the ES calculation are placed in the non-modelable risk bucket and subjected to stress tests that are similar to the GMS stress tests. The non-modelable risk factor add-on is very significant, expected to contribute 34.1 percent of the market risk capital requirement for Group 1 banks when the FRTB is implemented.[10] However, this increase in market risk capital is not caused by an increase in risk, but rather because the FRTB methodology repeats the GMS stress test for non-modelable risk factors. This problem could be best addressed through a general update to the GMS methodology, which will be covered in a future post.

Volatility of market risk capital

To use the ES methodology, a bank must show that it can satisfy two P&L attribution tests, described further below, that measure the degree to which risk factors used in the bank's risk management systems reflect the key risk drivers of the firm's P&L, so that oversimplifications in a bank's risk systems can be detected. Although these tests are sensible conceptually, in practice they are very difficult to implement. Often, the tests fail not because of simplifications in the risk systems, but because of data quality errors, which can cause an abrupt switch from the models-based calculation of market risk capital to the much larger standardized method.

The P&L attribution tests do not use actual P&L but rather compare Hypothetical P&L (HPL) with Risk-theoretical P&L (RTPL). HPL is calculated by valuing positions held at the end of the previous day using risk factors observed at the end of the subsequent day. RTPL, on the other hand, values the positions using only the risk factors that are

[10] See "Basel III Monitoring Report," September 2022

included in the firm's ES calculation, including the non-modeled risk factors. By comparing HPL to RTPL, the P&L attribution tests aim to detect when the set of risk factors used in ES along with the non-modeled risk factors are insufficient in number to capture the bank's risk. The P&L attribution tests require two statistical tests, one that measures the correlation between HPL and RTPL and the other that measures the similarity between their empirical frequencies of occurrence. If a P&L attribution test is in the amber zone, banks will be assessed a capital surcharge. If a P&L test is in the red zone, then it must switch to using the standardized approach. Banks may remove the capital surcharge or return to using the modeled approach when they pass the P&L attribution tests.

However, banks may fail the P&L attribution tests not because the number of risk factors they are using is insufficient but rather because of data quality errors in the tests. Thus, in practice, market risk capital could be very volatile even when the bank is employing an effective set of risk factors in its ES and non-modeled risk factor calculations. This problem could be resolved by allowing banks more time to resolve data quality issues before having to include a capital surcharge or switch to the standardized approach and/or by relaxing the P&L attribution test thresholds.

Incentives to lower portfolio diversification

The ES calculation is justified in assuming that portfolio diversification may break down in a stressed market environment. However, not providing capital benefits for risk diversification also creates the incentive to reduce diversification, which is contrary to the intent of any regulatory capital regime. By assigning equal weight to full diversification of all assets and diversification only within an asset class, the ES tends to exaggerate the breakdown of correlation. It is plausible to assume that correlations may move in the same direction for equities and credit, but less plausible to suppose that FX rates and interest rates also move in the same directions, since FX and interest rates are determined by macro factors while equities and credit are heavily influenced by individual company risks. To better balance incentives to diversify risk with penalties for de-correlation, the weighting scheme could be easily changed. For example, the diversified ES calculation discussed in Problem 4 could have been weighted 60 percent or 70 percent and the non-diversified ES could have been weighted 40 percent or 30 percent, respectively.

Conclusion

By working through an example in detail, we can see how the new FRTB ES-based capital measure addresses the five problems implicit in market risk capital based on VaR. The FRTB ES methodology may seem mysterious at first glance, but it is really nothing more than a generalization of techniques that have already been applied in the Basel II.5 standard. Although the FRTB fixes the problems with VaR, it also introduces new problems that should be addressed before FRTB implementation. These problems fortunately have simple solutions that do not require any significant changes to the FRTB.

In assessing how the GMS can be updated, it is important to deconstruct the FRTB to assess what risks it is already covering so that the GMS does not account for the same risks. A future post will cover how the GMS can be updated.
