JUAN CÁRDENAS

Proposal and Comment Information

Title: Regulatory Capital Rule: Category I and II Banking Organizations, Banking Organizations with Significant Trading Activity, and Optional Adoption for Other Banking Organizations, R-1887

Comment ID: FR-2026-0007-01-C14

Submitter Information

Name: Juan Cárdenas

Submitted Date: 05/26/2026

The use of Expected Shortfall for regulatory capital would be a serious mistake, making the U.S. financial system less safe, and would not be an improvement over the current Market Risk Rule. Please see the attached file for further details.

Comments on Regulatory Capital Rule: Category I and II Banking Organizations, Banking Organizations with Significant Trading Activity, and Optional Adoption for Other Banking Organizations, R-1887

I am a former bank examiner, commissioned (safety and soundness) by the Federal Reserve, and hold a Ph.D. in mathematics. I spent several years as a practitioner, designing and implementing market risk models, as well as many years as a supervisor, reviewing and evaluating such models for regulatory capital.

These comments concern the proposed measure for Market Risk, and specifically the use of Expected Shortfall (ES). In brief: using ES for bank, or bank holding company, regulatory capital would be a serious blunder, for various conceptual and practical reasons, as explained below. To adopt supervisory language, attempting to use ES to calculate regulatory capital would be ***unsafe and unsound***.

[In what follows, "bank" will be used generically for banking organization. "Losses" will refer to trading losses.]

Objection 1. Regulatory capital predicated on ES, as a (probability-weighted) *average* of potential losses beyond a given percentile, is **conceptually unsound** because solvency is *not* defined or determined in an average sense. While at a theoretical level, the excesses and shortfalls of the potential losses, with respect to ES, offset each other, this is clearly impossible in a real-world situation. I.e., if a bank experiences a loss larger than ES, it is not possible to offset such a loss with smaller losses the bank could have theoretically experienced instead.

To illustrate with a stylized example, suppose a bank (under stress conditions perhaps) is about to face a 2.5% tail of the P&L distribution with the following possible outcomes and probabilities:

- $100MM loss (60% probability);
- $250MM loss (30% probability);
- $650MM loss (10% probability),

therefore ES = $200MM.

If bank stakeholders expected the bank's capital to cover all but the most extreme case, it would be preposterous to think ES would be the capital level they would require, since they would evidently require capital of at least $250MM. It would be risible for the bank (or regulators, should ES be adopted for regulatory capital) to claim that capital of $200MM perfectly protects the bank, arguing that if one were to average all the potential losses, then they would have been fine!

Similarly, even assuming stationarity of the P&L distribution, it would make little sense to think that losses over different time periods could offset each other. Again, conceptually, if in a given time period a bank experiences a loss smaller than ES, it is not required to store the "excess" capital for use in a subsequent period with a loss greater than ES. The reverse situation would be even more problematic: if the initial loss wipes out ES, one would hardly expect that the bank could be allowed to cover that loss with hypothetical future capital excesses over ES; indeed, who and what mechanism would finance that?

NB: The above objection is of course the statistical flaw of averages.

Objection 2. Calculation of ES is *not* a well-posed problem. While ill-posed problems are of interest in mathematics, it would be highly imprudent to deliberately choose an ill-posed problem to be the basis for regulatory capital, given its vital role in the safety of the financial system.

In mathematics, a well-posed problem must satisfy three conditions:

A. The solution exists;
B. The solution is unique; and
C. The solution depends continuously on initial data/parameters[1].

ES fails the first and third requirements, which, as shown below, will have serious and undesirable practical implications.

Lack of existence: There are P&L probability distributions for which ES does not exist. Notably, this occurs, for example, in the case of a Cauchy distribution. More generally, for continuous distributions, ES fails to exist whenever the probability density $f(x)$ exhibits decay at infinity as:

$$f(x) \sim \frac{1}{x^{1+\alpha}}$$

for $0 < \alpha \leq 1$.

Analogously, ES can fail to exist for discrete probability distributions, e.g., let the probability $p(n)$ for a profit or loss of n, be

$$p(n) = \frac{3}{\pi^2 n^2} \text{ for } n = \pm 1, \pm 2, \pm 3, \cdots$$

[1] For suitably defined spaces and norms.

Although a legitimate probability distribution, as

$$\sum_{\substack{-\infty \\ n\neq 0}}^{\infty} p(n) = 1,$$

ES does not exist for this distribution (since the harmonic series diverges). Indeed, attempting to "estimate" ES by drawing N random samples (e.g., mimicking a simulation-based approach) and averaging the worst m outcomes (where $m / N \approx q$, the desired percentile, e.g., q = 2.5%) will fail and the "answer" will not be meaningful. The reason is that there is nothing to converge to in this case. In fact, any attempt at estimating ES in such cases will yield random results.

As an example, let us attempt to replicate a 4-year historical simulation approach to "estimate" ES in the case of the above distribution $p(n)$. In a simple experiment, running 10 trials, in which each trial draws 1,000 random samples, and then calculating the average of the worst 25 outcomes produced "answers" that ranged from -41.4 to -218.8. Note that the 99% confidence level VaR is -30 in this case.

Remark 1: under stressed market conditions, when capital is most needed, bank P&L distributions will be more likely to exhibit fatter tails, and hence more likely for ES to not exist.

Remark 2: Stable (non-normal) distributions, which have undefined mean (and hence ES will not exist) for a wide range of parameter values, may model financial data better, and best fit tail quantiles of empirical financial data distributions.[2]

Let us address possible counterarguments to the lack of existence for ES:

I. *This is only observed because a simulation approach (historical or Monte-Carlo) only draws a finite number of samples; if one were to draw more samples, then the ES estimate would converge.*

   This is not correct. For distributions where ES does not exist, drawing more samples will simply result in the "estimate" continuing to take on random values without ever converging.

II. *We are not worried about such cases because banks do not have infinite losses.*

   This is a misunderstanding of the reason for ES not converging. While banks' losses may be finite, this is not sufficient to ensure convergence. The reason ES

---

[2] See for example: Mandelbrot, B., New methods in statistical economics. *The Journal of Political Economy*, 71 #5, 421-440 (1963), and S. Mittnik, T. Rachev, and M.S. Paolella. Stable Paretian modelling in Finance: Some empirical and theoretical aspects. In Adler et al, editor, *A Practical Guide to Heavy Tails*. Birkhauser, 1998.

diverges is that as the losses become larger, their corresponding probabilities do not decay (i.e., go to zero) fast enough.

III. *Since we are defining ES as a measure of the average of all potential losses exceeding the VaR at a given confidence level, there is no ambiguity in calculating an average (possibly weighted) of extreme losses.*

This is a mistake because it confuses the algorithm to estimate ES terminating, with the estimate actually converging to something meaningful (which logically cannot happen when the quantity in question does not exist).

IV. *The standardized method would be available for situations where ES is not defined.*

There is currently no provision or expectation in the proposed rule for regulators or banks to test whether their P&L distributions have well-defined ES or not. Even if such an expectation were added to the rule, any attempt to implement such a test would be doomed because, as will be shown below, two P&L distributions can differ by an arbitrarily small amount, yet one could have a well-defined ES, and the other would not.

Thus, it would be **unsafe** to insist on calculating regulatory capital using ES when it is not defined, as it would generate meaningless, randomly varying results.

Lack of continuous dependence on initial data/parameters: small variations in the inputs to the ES calculation can have a large impact on the result (i.e., on ES). This is a serious concern because there will always be small measurement errors in the inputs (e.g., valuation models are approximations, inaccuracies in simulated market data, etc.), but these small errors can produce large errors in the predicted ES, rendering it untrustworthy.

1. ES can exhibit a large change with a small variation in the distribution.

   To illustrate with an example, suppose that in the 2.5% loss tail we have:

   - Loss = $10 in 99 out of 100 outcomes.
   - Loss = $99,010 in 1 out of 100 outcomes.
   - Then ES = $1,000

   If the loss distribution shifts slightly, the 1 in 100 loss is now $199,010, then ES doubles to $2,000. Calculating the L^∞ norm of the difference between the two distributions, we obtain 0.025% (= .01 × 2.5%).[3] Obvious modifications could make the difference as small as desired, yet with large impact on ES.

[3] Other norms would also measure small differences in the distributions.

2. A small change in the distribution can result in the ES going from well-defined to not defined. This is easily seen with the following example: suppose the density of the P&L distribution decays as

$$f(x) \sim \frac{1}{x^{2+\varepsilon}}$$

for an arbitrarily small $\varepsilon > 0$, then ES is well-defined. If instead $\varepsilon = 0$, then ES does not exist, as seen before. Note that this means that any attempt to design a test to detect whether a P&L distribution has changed from having a well-defined ES to one that does not, will fail because the distribution can vary by an amount smaller than the sensitivity of the test.

Other flaws of ES

ES is inefficient. ES requires capital for potential extreme losses that one would never intend to cover. To illustrate, return to the example where in the 2.5% loss tail we have:

- Loss = $10 in 99 out of 100 outcomes.
- Loss = $99,010 in 1 out of 100 outcomes.
- So ES = $1,000

If the loss to be realistically covered is $10, why require capital 100 times larger? Clearly, the excess $990 in capital has no hope of covering the extreme case (loss of $99,010).

While a possible counterargument might be that this extra capital is a charge or deterrent against accumulating tail risk, is it credible to suggest that (in this example) 1% capital is really a deterrent? Clearly, loss probabilities and amounts can be adjusted such that the capitalization ratio is as small as desired. This leads to the next flaw.

ES can be "gamed". A trader can engineer his/her P&L profile such that the 2.5% loss tail has the probability of a small loss is N times that of an extreme loss, and then the capital charge (ES) would be approximately $1/N$. Thus, capital based on ES would do little to deter the sale of deep out of the money options to collect premium.

To illustrate how the ratio of ES to a large loss M can be made arbitrarily small, let k be given; set:

- $P(loss = 1) = \frac{10^{2k}-1}{10^{2k}}$
- $P(loss = M) = \frac{1}{10^{2k}}$
- $M = 10^k + 1$

- then $\frac{ES}{M} = \frac{1}{10^k}$

ES fails to differentiate tail risk. Contrary to the claim that *"…an expected shortfall-based measure better reflects the tail risk of extreme events,…"* it is possible to have dramatically different tail risk profiles all with the same ES. For example, if, conditional on being in the 2.5% tail of the distribution, we have the following four different risk profiles:

- Case 1: $P(loss = 20)$ = 100%
- Case 2: $P(loss = 10) = P(loss = 30)$ = 50%
- Case 3: $P(loss = 1+2k)$ = 0.25% for k = 0…19
- Case 4: $P(loss = 10)$ = .6; $P(loss = 25)$ = .3; $P(loss = 65)$ = .1

Yet in all cases, the expected shortfall is 20, so ES has failed to give a better insight into the tail risk.

The above four cases also illustrate another deficiency of ES:

Solvency (in a VaR sense) provided by ES is not stable. We can see that solvency varies while ES stays constant. Specifically:

- Case 1: Solvency provided by ES = 20 covers 100% of losses.
- Case 2: Solvency is 98.75%
- Case 3: Solvency is 99.9375%
- Case 4: Solvency is 99%

It would not seem to be an improvement over VaR if the confidence level provided by an ES-based capital requirement were to vary as ES remains constant.

ES would suffer from large estimation error (for simulation approaches). I do not elaborate on this point, since this deficiency has already been covered by Danielsson and Zhou (2016) in *Why is risk so hard to measure.*[4]

Some key points from that paper:

- ES is estimated with more uncertainty than VaR (even in ES at 97.5% and VaR at 99% levels)
- It would require 50 years of daily data for estimators to achieve asymptotic properties
- Large uncertainty at smaller (<500) sample sizes

---

[4] Danielsson, Jon and Zhou, Chen, Why Risk Is So Hard to Measure (June 22, 2016). De Nederlandsche Bank Working Paper No. 494, Available at SSRN: https://ssrn.com/abstract=2597563 or http://dx.doi.org/10.2139/ssrn.2597563

- Upper bound of 99% confidence level bands are multiples of the forecast
- Note that a 1-year historical simulation would have to estimate ES(97.5%) by averaging 6 observations, and a 3-year historical simulation would only employ 19 observations to calculate a tail average.