Proceedings of the 2022 Winter Simulation Conference
B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C.G. Corlu, L.H. Lee, E.P. Chew, T. Roeder, and P. Lendermann, eds.

SEQUENTIAL NESTED SIMULATION FOR ESTIMATING EXPECTED SHORTFALL

Ou Dang

Insurance Risk and Finance Research Centre
Nanyang Business School
Nanyang Technological University
50 Nanyang Avenue
Singapore, 639798, SINGAPORE

Ben Feng

Department of Statistics and Actuarial Science
University of Waterloo
200 University Ave W
Waterloo, ON N2L 3G1, CANADA

ABSTRACT

Expected shortfall (ES) is a widely used tail risk measure in the financial industry. Estimating the ES of a financial portfolio usually requires nested simulation, which is computationally burdensome. In a standard nested simulation procedure, one first simulates a set of plausible evolution of underlying risk factors, or the scenarios. Then, conditional on each outer scenarios, inner simulations are run to evaluate the financial positions in that scenario. In this work, we propose a sequential nested simulation procedure that dynamically allocates a fixed simulation budget to accurately estimate the expected shortfall. The goal is to gradually concentrate the simulation budget on the tail scenarios with the largest losses, as these scenarios are most relevant in ES estimation. Our numerical experiments show that, with the same simulation budget, the proposed procedure significantly improves the estimation accuracy of ES compared to a standard nested simulation procedure.

1 INTRODUCTION

Estimating the risks, especially tail risks, of financial or insurance portfolios is an important and ubiquitous enterprise risk management task in the financial and insurance industry. Common tail risk measures include the probability of large loss, Value-at-Risk (VaR), Expected Shortfall (ES, also known as the Conditional Value-at-Risk or CVaR), etc.. These tail risk measures help companies understand their risk positions, which then informs further risk management decisions such as re-balancing the financial portfolios or setting aside appropriate capital reserves to maintain the overall financial solvency. We focus on the ES in this work because it is adopted by banking and insurance regulations such as BASEL III (Basel Committee on Banking Supervision 2019), NAIC Valuation Manual (NAIC 2022), and Life Insurance Capital Adequacy Test (OSFI 2019).

Consider the following example: Let the current time be time-0 and the longest maturity of the constituent assets in a portfolio of interest be $T > 0$; e.g., $T = 1$ year. A portfolio manager may be interested in monitoring the tail risk of the portfolio's profit and loss (P&L) at a given risk horizon τ where $0 < \tau < T$; e.g., $\tau = 1$ month. At any time $0 < t < T$, let V_t be the portfolio's fair market value (discounted to time-0) and so $V_0 - V_t$ is the portfolio P&L at time-t. For notational convenience, we define the portfolio loss at time-τ as $L = V_0 - V_\tau$. A risk measure $\rho(L)$ is a functional of L. For example, given a confidence level α, e.g., $\alpha = 95\%$, the α-ES is defined as $\alpha\text{-ES}(L) := E[L|L \geq \mathrm{VaR}_\alpha]$, where $\mathrm{VaR}_\alpha(L) := \sup\{l : P[L \geq l] > \alpha\}$ (Yamai and Yoshiba 2005). For a given sample of losses $L_1, \ldots, L_M$ and assume αM is an integer for simplicity, the α-ES is given by

$$ES_\alpha = \frac{1}{(1-\alpha)M} \sum_{i=\alpha M+1}^{M} L_{[i]}, \tag{1}$$

where $L_{[1]} \leq L_{[2]} \leq \cdots \leq L_{[M]}$ are the order statistics of the sample L_i, $i = 1, \ldots, M$. This is the quantity that we are interested in estimating in this work.

Nested simulation is a common tool to estimate tail risks. Typically, the portfolio value V_t depends some underlying risk factors, denoted by S_t, such as stock prices, bond prices, interest rates, etc. While the current values S_0 and V_0 maybe known or observed, the portfolio value at the future risk horizon, V_τ, is a random variable due to the stochasticity of S_τ. In a standard nested simulation procedure, the outer-level simulation generates, under the real-world measure, M plausible realizations of the risk factor S_τ, i.e., S^i_τ, $i = 1, \ldots, M$, which are known as the outer scenarios. In this study, we assume that the fair portfolio value V_t can be valuated as the discounted payoffs under the risk-neutral measure (e.g., the $\mathbb{Q}$-measure). That is, V_τ can be written as a conditional expectation given S_τ, i.e., $V_\tau = V(S_\tau) = E^{\mathbb{Q}}[H(S_{\tau^+})|S_\tau]$, where $H(S_{\tau^+})$ is the discounted payoffs given S_τ. In practical risk management applications, the dynamic asset model for S_t and the discounted payoffs $H(\cdot)$ are often complex so the conditional expectation $E^{\mathbb{Q}}[H(S_{\tau^+})|S_\tau]$ could not be calculated analytically. In such cases, given an outer scenario S^i_τ, $i = 1, \ldots, M$, the portfolio value in that scenario, i.e., $V^i_\tau = V(S^i_\tau) = E^{\mathbb{Q}}[H(S_{\tau^+})|S^i_\tau]$, is estimated via an inner-level simulation (or inner simulation for short). Specifically, in a standard nested simulation procedure, V^i_τ is estimated by $\widehat{V}^i_\tau = \frac{1}{N}\sum_{j=1}^N H(S^{i,j}_{\tau^+})$, where the inner replications $S^{i,j}_{\tau^+}$, $j = 1, \ldots, N$, are i.i.d. realizations of the conditional random variable $S_{\tau^+}|S^i_\tau$; the corresponding portfolio P&L is then $\widehat{L}_i = V_0 - \widehat{V}^i_\tau$. After running the inner simulations for all M scenarios, α-ES is then estimated by

$$\widehat{\alpha\text{-ES}} = \frac{1}{(1-\alpha)M} \sum_{i=\alpha M+1}^{M} \widehat{L}_{[i]} = \frac{1}{(1-\alpha)M} \sum_{i=\alpha M+1}^{M} \frac{1}{N} \sum_{j=1}^{N} \left(V_0 - H(S^{i,j}_{\tau^+}) \right), \tag{2}$$

where $\widehat{L}_{[1]} \leq \widehat{L}_{[2]} \leq \cdots \leq \widehat{L}_{[M]}$ are the order statistics of the portfolio P&L's $\widehat{L}_i$, $i = 1, \ldots, M$. We refer to the scenarios whose portfolio P&L's are averaged in (2) as the *tail scenarios* and the other scenarios as the *non-tail scenarios*.

For a given set of outer scenarios $S^1_\tau, \ldots, S^M_\tau$, one difference between (1) and (2) lies in the portfolio P&L L_i versus $\widehat{L}_i$ for all $i = 1, \ldots, M$: L_i are calculated in closed-form in (1) but are estimated via inner simulation in (2). Due to randomness in the inner simulation, the portfolio P&L's in (2) are subject to estimation errors, which can lead to a set of tail scenarios that is different from the the true tail scenarios in (1).

Though the randomness in the outer simulation is also a source of estimation error for the α-ES estimate, for ease of exposure we focus on designing an efficient simulation procedure that reduces the inner simulation estimation error for a given set of outer scenarios. Methods to reduce the noise from outer scenario simulation is outside the scope of this paper. Interested readers may refer to work such as Gordy and Juneja (2010) and Liu and Staum (2010) for more detailed discussion on how the number of out scenarios should be determined in a nested simulation. In general, if the number of out scenarios is too small, the variance of a nested simulation estimator tend to increase.

In practice, a financial portfolio can consist of thousands or tens of thousands of constituent contracts, whose payoffs may be complex. Moreover, many insurance portfolios have maturities that are a few decades long and the contract payoffs are path-dependent (Dang 2021). Thus, a standard nested simulation procedure as described above can be computationally demanding, sometimes prohibitively so. We observe in (2) that only the (estimated) tail scenarios' P&L's are averaged in estimating ES, the P&L's in all other scenarios were needed only for ranking the losses. Allocating the same computations to all scenarios, like in standard nested simulation, is a wasteful use of simulation budget. Intuitively, if one could efficiently and confidently differentiate tail scenarios from non-tail scenarios, then one should concentrate simulation budget to the tail scenarios. This intuition is the main inspiration for this work.

1.1 Literature Review

Despite the computational burden, nested simulation remains a useful tool in the risk management and financial reporting of financial portfolios. Thus, much research have been devoted to reducing the computational burden or using a given simulation budget efficiently.

Broadly speaking, existing literature addresses the computational burden from two different angles: The first is to replace the Monte Carlo inner simulation with other methods such as regression (Broadie et al. 2015; Jiang et al. 2020), kernel smoothing (Hong et al. 2017), stochastic kriging or Gaussian Process regression (Liu and Staum 2010; Risk and Ludkovski 2018), likelihood-ratio weighted re-sampling (Feng and Staum 2017), and neural network (Fernandez-Arjona 2021). In regression-based methods, choosing appropriate basis functions is both critical and challenging. In kriging-based methods, picking design points can be difficult in some applications. In the likelihood-ratio weighted re-sampling method, the likelihood calculation overhead can be significant if the experiment is not carefully designed (Dang 2021).

The second angle is to strategically allocate a fixed computation budget to minimize the estimation error, which is typically measured by the mean squared error (MSE). The strategic allocation can be the optimal numbers of outer scenarios and inner simulations in a standard nested simulation (Gordy and Juneja 2010; Zhang et al. 2022). Specifically, the given simulation budget is uniformly allocated accross all scenarios. However, in practice, sometimes the outer scenarios are provided by exogenous sources such as the modeling team in a company or the regulator (Risk and Ludkovski 2018) so the users cannot choose or alter the scenarios. Budget allocation can also be non-uniform for different scenarios. For example, Lan et al. (2010) propose a two-stage procedure to estimate the ES and constructs a confidence interval around it: A low-budget initial inner simulation is conducted in the first stage to identify the tail scenarios via hypothesis test. In the second stage, the remaining computation budget is allocated uniformly to the outer scenarios identified in the first stage. A confidence interval for the ES is constructed using the simulation output from the second stage. Broadie et al. (2011) propose a sequential simulation procedure to estimate the probability of large loss over a known threshold. Their proposal also starts with a low-budget initial inner simulation then iterates sequentially. In each subsequent iteration, an additional inner simulation is allocated to one outer scenario that is most likely to be around the large loss threshold. The procedure then repeats until the given simulation budget is exhausted. In Giles and Haji-Ali (2019) and Alfonsi et al. (2021), a multi-level Monte Calro procedure is proposed. At each level, the number of outer scenarios is a function of some error tolerance and the number of inner simulations in a given outer scenario is related to the ratio of absolute value of sample mean over sample standard deviation. Bouchard et al. (2021) presents a multi-step procedure where in each step the inner simulation progressively shift towards a smaller subset of tail outer scenarios that are relevant to the ES estimation. The optimal number of inner simulations and tail outer scenarios in each step can be solved by dynamic programming and neural network approximation.

Our work is inspired by the aforementioned research, particularly by non-uniform budget allocations with concentrations towards tail scenarios such as Broadie et al. (2011), Lan et al. (2010), and Bouchard et al. (2021). We identify two main drawbacks in the proposal by Lan et al. (2010): Firstly, if the variance of the inner-level simulation is large, the first stage may require a significant computation budget to properly screen out a meaningful portion of scenarios. Secondly, when a large set of outer scenarios are given, the hypothesis test in the first stage requires significant computation. The threshold for large loss is known in Broadie et al. (2011), which simplifies the tail versus non-tail categorization as computations are concentrated only to the region surrounding the known threshold. However, when estimating ES, the loss threshold for the tail scenario, i.e., the value-at-risk, also needs to be estimated. So the methods in Broadie et al. (2011) cannot be directly extended to estimating ES. In Bouchard, Reghai, and Virrion (2021), identifying the optimal number of inner simulation and tail scenarios in each adaptive step is complex and computationally costly, although the authors argue that such computation can be done off-line.

The main contribution of this work is proposing and testing an efficient sequential simulation procedure that strategically allocates a fixed simulation budget to accurately estimate ES. Specifically, through sequential iterations, the proposed procedure identifies and updates tail scenarios towards which computations are gradually concentrated. It is intuitive and easy to implement. Moreover, the proposed procedure requires no prior knowledge of the portfolio's payoff structure thus is applicable to a wide range of financial and insurance risk management applications.

This paper is organized as follows: In Section 2, we outline the proposed procedure and provide detailed discussions on the rationale and the choice of design parameters. In Section 3, we present numerical examples to demonstrate the the effectiveness of the proposed procedure. In Section 4 we conclude this paper and discuss future work.

2 SEQUENTIAL NESTED SIMULATION OF EXPECTED SHORTFALL

In this section, we present a two-stage sequential nested simulation procedure for estimating the α-ES in (1) for a given set of scenarios $S_\tau^1, \ldots, S_\tau^M$ and a fixed simulation budget, i.e., a fixed total number of inner replications. Section 2.1 provides an intuitive high-level description of the procedure. Section 2.2 presents the algorithm and detailed discussions on the choice of experiment design parameters.

2.1 Estimating Expected Shortfall via Nested Simulation

As discussed in the introduction, estimating α-ES using the standard nested simulation given in (2) is a wasteful use of simulation budget: In risk management applications, the confidence level α is typically high, e.g., $\alpha = 95\%$ or $\alpha = 99\%$. In these cases, there are only few tail scenarios whose estimated losses are averaged in the ES estimation. The non-tail scenarios' losses provide a ranking of all the losses, but do not directly impact the value of the ES estimate. This suggests that we should screen out non-tail scenarios with as little simulation budget as possible then concentrate the remaining budget to the tail scenarios. This screening process can only be done with some pilot experiments because we can only tell if a scenario belongs to the tail or not after running some inner simulations.

Based on the above observations, we propose a two-stage sequential nested simulation procedure. Stage 1 includes an iterative process where small-scale inner simulations are run. The small-scale pilot simulations are then used to form confidence intervals for portfolio P&L estimates in different scenarios, some of which are then identified as non-tail scenarios and are screened out from Stage 2. Stage 1 aims to obtain reasonable loss estimates in all scenarios in an efficient manner so that the scenarios in or near the tail region of the loss distribution can be properly identified for Stage 2. Stage 2 sequentially allocates the remaining simulation budget within this set scenarios. Specifically, additional simulation budgets are allocated to scenarios with the largest losses. In each sequential iteration in Stage 2, the portfolio P&L estimates and sample variances are updated, which then lead to an updated set of tail scenarios. Stage 2 aims to concentrate the remaining simulation budget only on the tail scenarios, which leads to an accurate α-ES estimate. We will illustrate in Section 3 that both stages serve their respective purposes well.

Stage 1 in our proposed procedure is inspired by the screening process in Lan et al. (2010). However, instead of conducting a stringent hypothesis test, we use the confidence intervals for individual portfolio P&L estimates to screen out scenarios. Our approach avoids the computationally heavy pairwise comparisons. The sequential allocation method in Stage 2 is inspired by Broadie et al. (2011). Rather than concentrating simulation budget near a known large loss threshold, our approach allocate simulation budget in all tail scenarios.

2.2 Algorithm

Algorithm 1 outlines the steps of the two-stage sequential nested simulation process, followed by more detailed explanation of each step.

In this two-stage nested simulation process, a total simulation budget of Γ is used. The user also designates part of the budget, i.e., Γ_1, to Stage 1. Both stages include multiple iterations of nested simulation. Each iteration uses a simulation budget of γ. The process is initialized by a standard nested simulation with a small and equal number of inner replications in all scenarios. In each subsequent iteration

Algorithm 1: Sequential nested simulation for estimating α-ES.

input : – Underlying real-world and risk-neutral models with parameters.
– M outer scenarios, each is a real-world risk factor sample, $S^i_\tau, i = 1, \ldots, M$.
– Γ: Total computation budget.
– Γ_1: Computation budget for Stage 1.
– m: Number of outer scenarios simulated in *each* iteration in Stage 2.
– γ: Computation budget for *each* iteration of inner simulation.
– β: Confidence level of the confidence interval for $\widehat{L}_i$ in Stage 1.

output: $\alpha\text{-}\widehat{\text{ES}}(L)$ for the portfolio of interest.

Initialization: Simulation set $\Omega \leftarrow \{S^i_\tau : i = 1, \ldots, M\}$; Remaining computation budget $\Theta \leftarrow \Gamma$; Inner simulation sample size $n_i \leftarrow \frac{\gamma}{|\Omega|}$, $i = 1, \ldots, M$; Cumulative inner simulation sample size $N_i \leftarrow 0$

Stage I: Pilot simulation

while $\Theta \geq \Gamma - \Gamma_1$ **do**

Nested simulation for outer scenarios $i \in \Omega$ with n_i inner simulation in each outer scenario.

$N_i \leftarrow N_i + n_i$; $\quad \Theta \leftarrow \Theta - \gamma$

Update sample mean $\widehat{L}_i$ and sample standard deviation s_i of L_i for $i \in \Omega$ with the latest simulation output.

Update confidence interval for each L_i for $i \in \Omega$. The confidence interval is

$$(LB_i, UB_i) \leftarrow \left(\widehat{L}_i - t_{(1-\beta)/2, N_i - 1} \frac{s_i}{\sqrt{N_i}}, \widehat{L}_i + t_{(1-\beta)/2, N_i - 1} \frac{s_i}{\sqrt{N_i}} \right)$$

.

$\Omega \leftarrow \{i : UB_i > LB_{[\alpha M+1]}, i = 1, \ldots M\}$ where $LB_{[i]}$ is the i-th order statistics of LB_i, for $i = 1, \ldots M$.

$n_i \leftarrow \max \left(\left(\gamma + \sum_{i \in \Omega} N_i \right) \frac{s_i^2}{\sum_{i \in \Omega} s_i^2} - N_i, 0 \right)$.

end

Stage II: Nested simulation with concentrated computation budget

$\Omega \leftarrow \{i : \widehat{L}_i > \widehat{L}_{[(M-m)]}, i = 1, \ldots M\}$.

while $\Theta \geq 0$ **do**

$n_i \leftarrow \max \left(\left(\gamma + \sum_{i \in \Omega} N_i \right) \frac{s_i}{\sum_{i \in \Omega} s_i} - N_i, 0 \right)$.

Nested simulation for outer scenarios $i \in \Omega$ with n_i inner simulation in each outer scenario.

$N_i \leftarrow N_i + n_i$; $\quad \Theta \leftarrow \Theta - \gamma$

Update sample mean $\widehat{L}_i$ and sample standard deviation s_i of L_i for $i \in \Omega$ with the latest simulation output.

$\Omega \leftarrow \{i : \widehat{L}_i > \widehat{L}_{[(M-m)]}, i = 1, \ldots M\}$.

end

Estimate the $\alpha\text{-}\widehat{\text{ES}}(L)$ as $\alpha\text{-}\widehat{\text{ES}}(L) = \frac{1}{(1-\alpha)M} \sum_{i=\alpha M+1}^{M} \widehat{L}_{[i]}$.

of Stage 1, the sample means, standard deviations, and confidence intervals of the portfolio P&L estimates in some scenarios are updated with new inner simulations conducted from the current iteration. In the next iteration, we screen out scenarios whose upper bound of the confidence interval is lower than the $(\alpha M+1)$-th highest lower bound. In other words, we screen out the outer scenarios that are unlikely to be tail scenarios based on the confidence intervals constructed so far. This process stops once Stage 1 uses up the designated simulation budget Γ_1.

In Stage 1, the number of inner replications n_i used in each iteration is calculated according to Line 6 of Algorithm 1 (rounded to the nearest integer). This calculation aims to make the surviving scenarios' confidence intervals in Line 4 have approximately equal widths in the next iteration, which reduces the overlap between confidence intervals and helps distinguish tail scenarios from the non-tail scenarios. To elaborate, let s_i' and N_i' denote the sample standard deviation and sample size for scenario i in the next iteration. A constant width of the confidence interval, i.e. $s_i'/\sqrt{N_i'}=k$ for some constant k and $i\in\Omega$, implies that $\sqrt{N_i'}\propto s_i'$, which is equivalent to $N_i'\propto s_i'^2$. Assuming $s_i=s_i'$ and given $N_i'=N_i+n_i$, to achieve a constant width of the confidence interval in the new iteration, we should then allocate the simulation budget γ so that $(N_i+n_i)\propto s_i^2$ for $i\in\Omega$, subject to a total cumulative simulation budget of $\gamma+\sum_{i\in\Omega}N_i$. Given N_i is non-decreasing, this leads to the allocation of n_i in Line 6 of Algorithm 1.

In each iteration of Stage 2, only the outer scenarios with the largest m estimated losses in the previous iteration receive additional inner simulations in the current iteration. After each iteration, the sample mean $\widehat{L}_i$ of the simulated outer scenarios are updated. In Stage 2, the number of inner simulation n_i used in each iteration is calculated according to Line 9 of Algorithm 1 (rounded to the nearest integer). This calculation aims at minimizing the variance of the $\alpha\text{-}\widehat{\text{ES}}$ estimate, subject to next iteration's simulation budget of γ. Therefore, the desired allocation of γ to each scenario $i\in\Omega$ satisfies the solution to this optimization problem:

$$\min_{N_i', i\in\Omega} \operatorname{Var}\left(\alpha\text{-}\widehat{\text{ES}}\right) \qquad \text{subject to} \qquad \sum_{i\in\Omega}N_i'=\gamma+\sum_{i\in\Omega}N_i. \tag{3}$$

Assuming the $\alpha\text{-}\widehat{\text{ES}}$ estimator averages over all the tail scenarios without any error in scenario ranking, then by Equation (2),

$$\operatorname{Var}\left(\alpha\text{-}\widehat{\text{ES}}\right)=\operatorname{Var}\left(\frac{1}{(1-\alpha)M}\sum_{i=\alpha M+1}^{M}\frac{1}{N_i'}\sum_{j=1}^{N_i'}\left(V_0-H(S_{\tau^+}^{i,j})\right)\right)=\frac{1}{(1-\alpha)^2M^2}\sum_{i=\alpha M+1}^{M}\frac{1}{N_i'}\operatorname{Var}\left(H(S_{\tau^+}^{i})\right).$$

Replacing $\operatorname{Var}\left(H(S_{\tau^+}^{i})\right)$ by the sample variance s_i^2, the optimization problem in (3) is reformulated as

$$\min_{N_i', i\in\Omega}\frac{s_i^2}{N_i'} \qquad \text{subject to } \sum_{i\in\Omega}N_i'=\gamma+\sum_{i\in\Omega}N_i \tag{4}$$

One can show that an optimal solution to (4) is $N_i'=N_i+n_i=\left(\gamma+\sum_{i\in\Omega}N_i\right)\frac{s_i}{\sum_{i\in\Omega}s_i}$. This leads to the allocation of n_i in Line 9 of Algorithm 1.

The rationale behind the Stage 1 simulation is to use a relatively small portion of the total simulation budget to screen out scenarios that are unlikely to be in the tail region. Stage 1 is efficient in that the scenarios that are far from the tail of the loss distribution is gradually eliminated by the comparison of the upper and lower bounds of the confidence interval for each outer scenario. The rationale behind the Stage 2 simulation is to concentrate the simulation budget only on a small set of tail scenarios to achieve higher accuracy in the α-ES estimate. The set of simulated scenarios in this stage can be small because the ranking of the scenario loss is already reasonably accurate after Stage 1. Note though in some applications, there are more efficient methods, such as proxy modeling in Dang, Feng, and Hardy (2020), to set apart tail versus non-tail scenarios.

As outlined in the input section of Algorithm 1, several design parameters need to be chosen for this process. The optimal choice of these parameters will be considered in future work. Below are some guidelines on how to choose the design parameters based on our experience:

- Γ_1, the simulation budget for the entire Stage 1 should be a small portion of Γ. In our experiment, we used $\Gamma_1 = 20\%\Gamma$. If Γ_1 is too small, it doesn't correctly identify the tail scenarios for Stage 2 simulation. If Γ_1 is too big, Stage 2 does not have sufficient remaining budget to concentrate on the tail scenario simulation. This will result in bias and lower accuracy of the α-ES estimate. In Section 3, we show in the numerical experiments the impact of using various values for Γ_1.
- m, the number of outer scenarios included in each iteration of Stage 2 should be larger than $(1-\alpha)M$, the number of tail scenarios in the Expected Shortfall calculation. This is to leave some margin for error in the ranking of scenarios identified in the iterative process. However, a large m increases the likelihood of wasting simulation budget on non-tail scenarios. In our experiments, we have chosen an m such that $P\left[l_{95\%} \geq \widehat{L}_{[M-m+1]}\right] = \sum_{i=M-m+1}^{M} \binom{m}{i} (1-\alpha)^i \alpha^{M-i} = 0.9999955$, where $l_{95\%}$ is the 95-th percentile of $\widehat{L}_i$. This is based on the assumption that $\widehat{L}_i$'s are i.i.d. Since each outer scenario uses different number of inner simulations, we cannot conclude that such assumption holds. Yet the result gives us some indication of how m should be chosen.
- γ, the simulation budget for each iteration, should be chosen in conjunction with the other design parameters to achieve a desired number of iterations in each stage. In general, smaller γ and more iterations will allow the simulation results get updated more frequently and improve the identification of tail scenarios. However, very small γ will result in many iterations and potentially high overhead calculation cost.
- β, the confidence level for the confidence interval of L_i, should be high for the confidence interval to be meaningful.

3 NUMERICAL EXPERIMENTS

In this section, we conduct numerical experiments using the following portfolio to illustrate the effectiveness of estimating an ES at $\alpha = 95\%$ level using our proposed sequential nested simulation process. The portfolio consists of three down-and-out put options with a term-to-maturity $T = 1/12$ and a risk horizon of $\tau = 1/52$:

- A long position with a strike price of 101 and a barrier of 91
- A long position with a strike price of 110 and a barrier of 100
- A short position with a strike price of 114.5 and a barrier of 104.5

We assume a current stock price of $S_0 = 100$. We also assume the stock price follows a geometric Brownian Motion with a drift of 8% and volatility of 20% under the real world measure. The risk free rate is 3%. Closed-form formula for the portfolio value V_t is available in this case (Haug 2007) so in theory no inner simulation is required. However, we use this example so that we can compare the simulation results with the true portfolio value.

The numerical experiments discussed in this section are conducted using one set of $M = 10^4$ outer scenarios and a total simulation budget of $\Gamma = 3\times 10^6$. The design parameters in the sequential nested simulation process are chosen to be: $m = 600$, $\gamma = 1.2\times 10^5 = 4\%\Gamma$, and $\beta = 95\%$. We have tested the impact of using different Γ_1 in the numerical experiments. We will clarify the Γ_1 value used when we present those results.

We first compare the loss distribution of the M outer scenarios from a sequential nested simulation with that from a standard nested simulation. Both simulation use the same simulation budget of Γ. In the sequential nested simulation, $\Gamma_1 = 6 \times 10^5 = 20\%\Gamma$. Figure 1 illustrates the output from these two experiments. We can see that with the same simulation budget, compared to the standard nested simulation, the simulated losses from the sequential simulation among tail scenarios, that is, the scenarios lying above the blue 95-th percentile line, have much less simulation noise and are much closer to the red dots representing the true loss of each scenario. This demonstrates the superior accuracy in estimating the tail loss of the sequential nested simulation process than the standard nested simulation process. Note though the higher accuracy in estimating the tail losses comes at the cost of lower accuracy in the non-tail scenario, as shown in the figure below the blue line. However, this has little impact on the accuracy of the α-ES estimate.

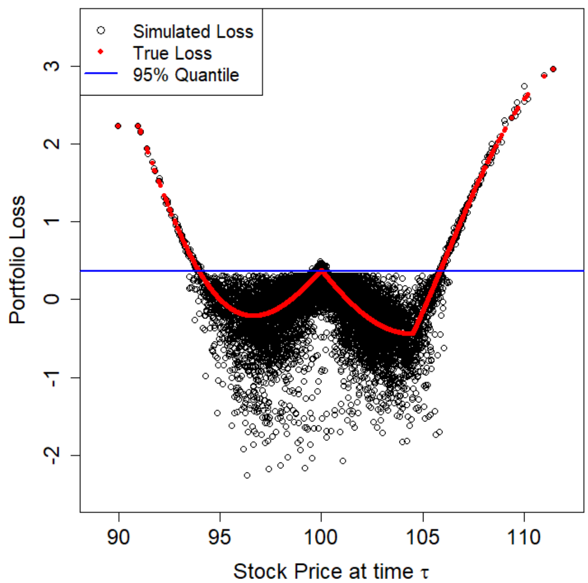

(a) Sequential nested simulation

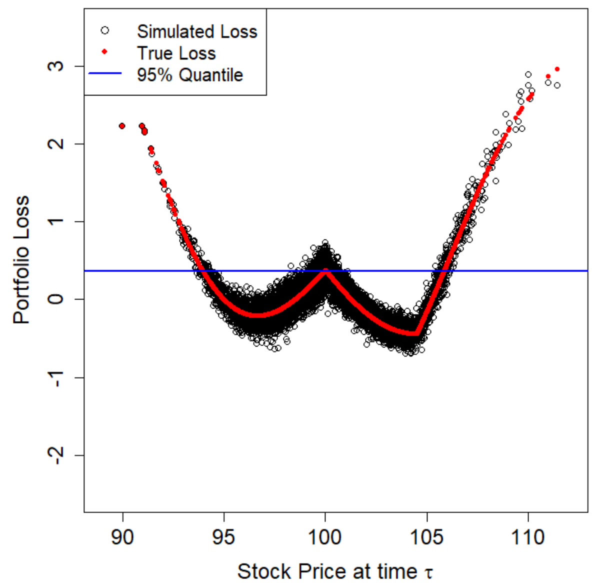

(b) Standard nested simulation

Figure 1: Loss distribution of a sequential nested simulation and of a standard nested simulation by S_τ^i, $i = 1, \ldots, M$. $M = 10,000$ outer scenarios. Black circles represent simulated loss. Red dots represent true loss calculated in closed-form. Blue line represent the 95-th percentile of the true loss distribution.

Figure 2 shows the number of inner simulations conducted in each outer scenario. The x-axis of the figure shows the ranking of each scenario in terms of true loss. Scenarios ranked higher than $9,500$ in this case are the tail scenarios. As illustrated in the figure, the tail scenarios have received the most inner simulations. The simulation budget is also more concentrated in scenarios that are close to the tail of the loss distribution.

Figure 3 compares how the 95%-ES estimate improves as more iterations are deployed in a sequential nested simulation experiment versus a standard nested simulation experiment. In the sequential experiment, the Stage 2 experiment starts after 2.4×10^6 remaining simulation budget is reached. The figure demonstrates that the 95%-ES estimate in the sequential experiment quickly improves in Stage 1 of the simulation, although at a slower speed than the standard nested simulation. The slower speed in the sequential experiment is due to the budget allocation strategy used in the Stage 1 simulation. It allocates less simulation budget to scenarios with a narrow confidence interval. As such, the α-ES estimate based on the Stage 1 output in a sequential experiment is less accurate than a standard nested simulation. Nevertheless, this does not reduce the accuracy of the α-ES estimate because more simulation budget is allocated to the relevant tail scenarios in Stage 2 to improve its accuracy. As shown in the figure, in the Stage 2 simulation, the 95%-ES estimate quickly reduces to the benchmark level, much faster than the standard nested simulation. In fact, with the same simulation budget, the 95%-ES estimated by the standard nested simulation never converges to the

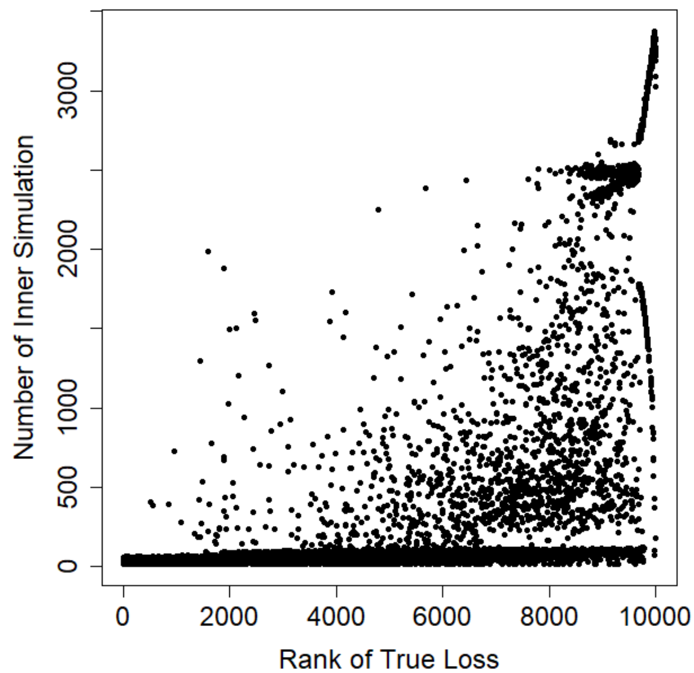

Figure 2: Number of inner simulations conducted by the rank of outer scenario.

benchmark 95%-ES. This observation is consistent with our intention for the sequential nested simulation process which was discussed in Section 2: Stage 1 is for setting apart tail and non-tail scenarios, in which the accuracy of the loss estimation itself is not critical, while Stage 2 is for improving the accuracy of tail loss estimation and reducing the variance of the α-ES estimate.

We also conduct 7 sets of sequential nested simulation and standard nested simulation experiments. We repeat each set of experiment 100 times and evaluate their accuracy. All the experiments use the same simulation budget Γ. For the sequential nested simulation, we conduct several sets of repeated experiment with different design parameters. Table 1 shows the detail of each experiment as well as the relative root mean squared error (relative RMSE) from each set of experiment. The relative RMSE is calculated as

$$\text{Relative RMSE} = \frac{\sqrt{\frac{\sum_{k=1}^{100}\left(95\%,k\text{-}\widehat{\text{ES}} - 95\%\text{-ES}\right)^2}{100}}}{95\%\text{-ES}},$$

which measures how much the RMSE deviates from the benchmark value of the estimate. Here 95%-ES represents the benchmark 95%-ES estimated by using the closed-form formula for V_t. The same results are also illustrated in the boxplot in Figure 4.

The results from these experiment show a clear advantage of the sequential nested simulation design, represented by Experiment (c), over other simulation designs. Given the same simulation budget, Experiment (c) achieves the smallest relative RMSE. From the boxplot, we can see that Experiment (c) has the smallest bias and very small variance among the eight sets of experiments. It is worth noting that all variations of the sequential nested simulation in Experiment (b) to (g) has smaller relative RMSE than the standard nested simulation. In Experiment (b), when no simulation budget is allocated to Stage 1 of the process, the 95%-ES estimate has a negative bias. This is driven by the fact that some tail scenarios are missing from the Stage 2 simulation of the experiments and the the ES estimation because Stage 1, which is meant for correctly identifying the true tail scenarios, is omitted. As a result, the ES estimation

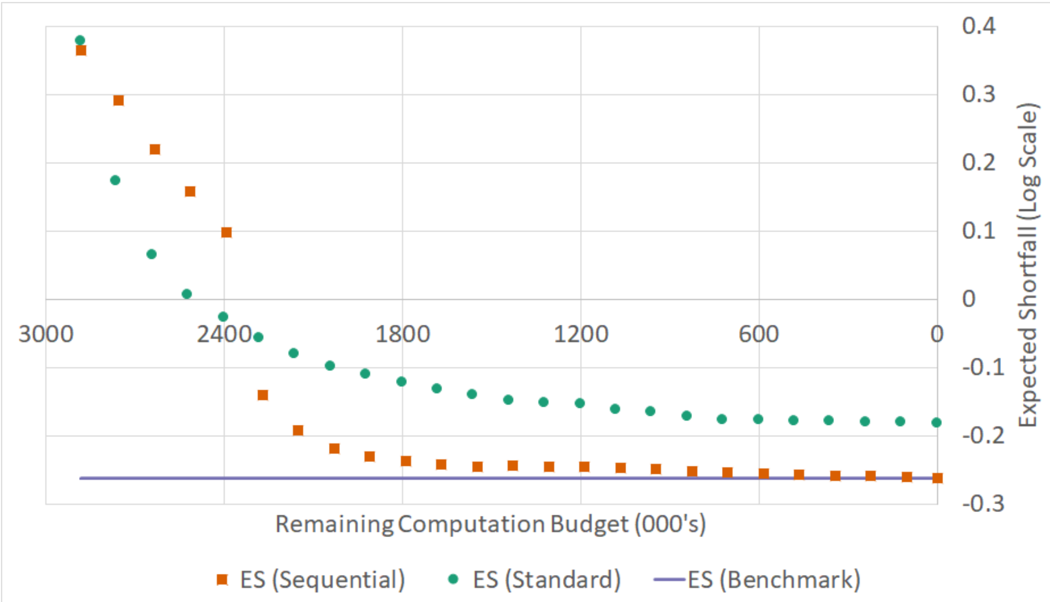

Figure 3: 95%-ES estimate by number of remaining simulation budget. x-axis is the remaining simulation budget in thousands.

Table 1: Relative root mean squared error (relative RMSE) from sequential and standard nested simulation.

	Experiment	Relative RMSE
(a)	Standard nested simulation with $N = 300$	8.825%
(b)	Sequential nested simulation with $\Gamma_1 = 0\%\Gamma$	5.022%
(c)	Sequential nested simulation with $\Gamma_1 = 20\%\Gamma$	0.361%
(d)	Sequential nested simulation with $\Gamma_1 = 40\%\Gamma$	1.081%
(e)	Sequential nested simulation with $\Gamma_1 = 60\%\Gamma$	1.680%
(f)	Sequential nested simulation with $\Gamma_1 = 80\%\Gamma$	2.524%
(g)	Sequential nested simulation with $\Gamma_1 = 100\%\Gamma$	6.226%

includes non-tail scenarios which inherently have smaller losses. In Experiment (d) to (g), when too much simulation budget is allocated to Stage 1 and not enough in Stage 2, the tail scenarios, even though being identified correctly, do not have an accurate loss estimate. This has the similar impact of having a standard nested simulation without sufficient inner simulations such as in Experiment (a) and is why a positive bias is observed.

4 CONCLUSION

In this paper, we present an efficient and easy-to-implement sequential nested simulation process for estimating Expected Shortfall. Stage 1 of the two-stage process is a budget-saving pilot simulation which deploys computation budget to the most relevant outer scenarios through a comparison of confidence intervals for the loss estimate. Stage 2 then sequentially allocate the computation budget to the tail scenarios and update the ES estimation in each iteration. The computation budget is allocated non-uniformly to each scenario to achieve higher accuracy of the α-ES. The proposed sequential nested simulation process is applicable to many asset and liability models without requiring strong assumptions. It can be widely adopted in finance and insurance applications for estimating Expected Shortfall.

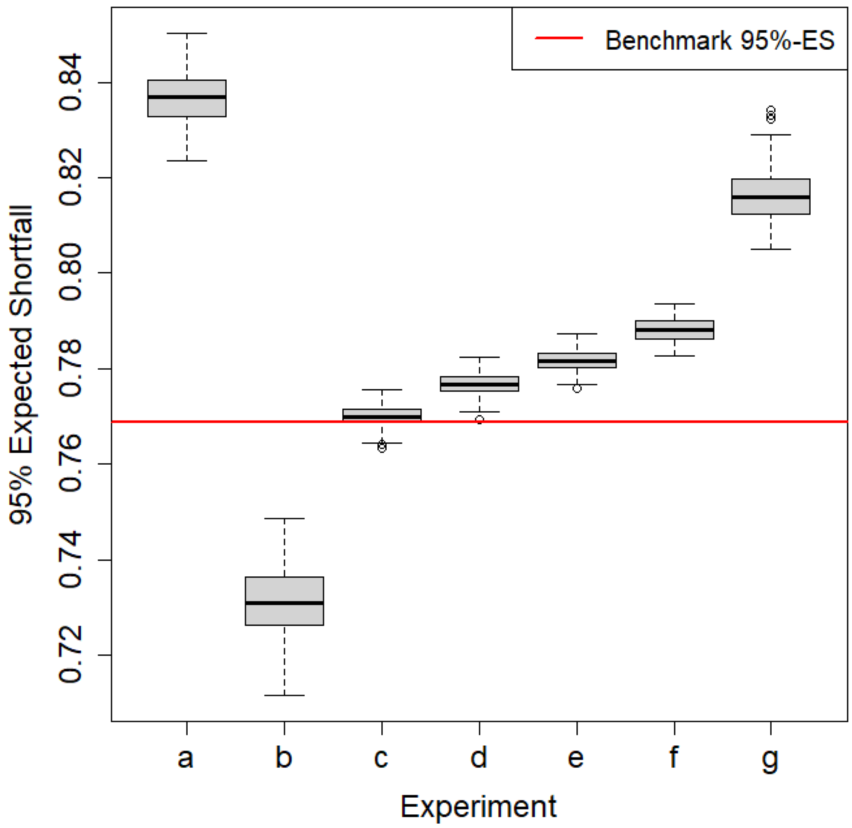

Figure 4: Relative RMSE of various nested simulation designs. (a) Standard nested simulation with $N = 300$; (b) Sequential nested simulation with $\Gamma_1 = 0\%\Gamma$; (c) Sequential nested simulation with $\Gamma_1 = 20\%\Gamma$; (d) Sequential nested simulation with $\Gamma_1 = 40\%\Gamma$; (e) Sequential nested simulation with $\Gamma_1 = 60\%\Gamma$; (f) Sequential nested simulation with $\Gamma_1 = 80\%\Gamma$; (g) Sequential nested simulation with $\Gamma_1 = 100\%\Gamma$.

For future work, we will study the convergence of the proposed process, the optimal selection of the design parameters, and extend it to estimating quantile risk measure such as Value-at-Risk. Conditional on capturing all the true tail scenarios in Stage 2, the convergence of our proposed process can be derived in a similar manner as Broadie, Du, and Moallemi (2011). Nevertheless, the probability of capturing all true tail scenarios in Stage 2 only approaches 1 when Stage 1 eliminates no outer scenario, and every outer scenario is considered in Stage 2.

ACKNOWLEDGMENTS

We acknowledge support from the Natural Sciences and Engineering Research Council of Canada (NSERC), funding reference number RGPIN-2018-03755.

REFERENCES

Alfonsi, A., A. Cherchali, and J. A. I. Acevedo. 2021. “Multilevel Monte-Carlo for Computing the SCR with the Standard Formula and Other Stress Tests”. *Insurance: Mathematics and Economics* 100:234–260.

Basel Committee on Banking Supervision 2019. “Minimum Capital Requirements for Market Risk”. Technical report, Basel Committee on Banking Supervision.

Bouchard, B., A. Reghai, and B. Virrion. 2021. “Computation of Expected Shortfall by Fast Detection of Worst Scenarios”. *Quantitative Finance* 21(7):1087–1108.

Broadie, M., Y. Du, and C. C. Moallemi. 2011. "Efficient Risk Estimation via Nested Sequential Simulation". *Management Science* 57(6):1172–1194.

Broadie, M., Y. Du, and C. C. Moallemi. 2015. "Risk Estimation via Regression". *Operations Research* 63(5):1077–1097.

Dang, O. 2021. *Efficient Nested Simulation of Tail Risk Measures for Variable Annuities*. Ph. D. thesis, Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Ontario. http://hdl.handle.net/10012/17084.

Dang, O., M. Feng, and M. R. Hardy. 2020. "Efficient Nested Simulation for Conditional Tail Expectation of Variable Annuities". *North American Actuarial Journal* 24(2):187–210.

Feng, M., and J. Staum. 2017. "Green Simulation: Reusing the Output of Repeated Experiments". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 27(4):1–28.

Fernandez-Arjona, L. 2021. "A Neural Network Model for Solvency Calculations in Life Insurance". *Annals of Actuarial Science* 15(2):259–275.

Giles, M. B., and A.-L. Haji-Ali. 2019. "Multilevel Nested Simulation for Efficient Risk Estimation". *SIAM/ASA Journal on Uncertainty Quantification* 7(2):497–525.

Gordy, M. B., and S. Juneja. 2010. "Nested Simulation in Portfolio Risk Measurement". *Management Science* 56(10):1833–1848.

Haug, E. G. 2007. *The Complete Guide to Option Pricing Formulas*, Volume 2. New York: McGraw-Hill.

Hong, L. J., S. Juneja, and G. Liu. 2017. "Kernel Smoothing for Nested Estimation with Application to Portfolio Risk Measurement". *Operations Research* 65(3):657–673.

Jiang, G., L. J. Hong, and B. L. Nelson. 2020. "Online Risk Monitoring Using Offline Simulation". *INFORMS Journal on Computing* 32(2):356–375.

Lan, H., B. L. Nelson, and J. Staum. 2010. "A Confidence Interval Procedure for Expected Shortfall Risk Measurement via Two-level Simulation". *Operations Research* 58(5):1481–1490.

Liu, M., and J. Staum. 2010. "Stochastic Kriging for Efficient Nested Simulation of Expected Shortfall". *Journal of Risk* 12(3):3.

NAIC 2022. "NAIC Valuation Manual". Technical report, National Association of Insurance Commissioners.

OSFI 2019. "Life Insurance Capital Adequacy Test". Technical report, Office of the Superintendent of Financial Institutes Canada.

Risk, J., and M. Ludkovski. 2018. "Sequential Design and Spatial Modeling for Portfolio Tail Risk Measurement". *SIAM Journal on Financial Mathematics* 9(4):1137–1174.

Yamai, Y., and T. Yoshiba. 2005. "Value-at-risk versus Expected Shortfall: A Practical Perspective". *Journal of Banking & Finance* 29(4):997–1015.

Zhang, K., G. Liu, and S. Wang. 2022. "Bootstrap-based Budget Allocation for Nested Simulation". *Operations Research* 70(2):1128–1142.

AUTHOR BIOGRAPHIES

OU DANG is a Research Assistant Professor in the Insurance Risk and Finance Research Centre at Nanyang Technological University. She holds a Ph.D. in Actuarial Science from the University of Waterloo. She is a Fellow of the Society of Actuaries (FSA) and a Fellow of the Canadian Institute of Actuaries (FCIA). Her research interest include simulation design, mortality modelling, life insurance product design and risk management. Her email address is ou.dang@ntu.edu.sg.

BEN FENG is an Assistant Professor in actuarial science at the University of Waterloo. He earned his Ph.D. in the Department of Industrial Engineering and Management Sciences at Northwestern University. He is an Associate of the Society of Actuaries (ASA). His research interests include stochastic simulation design and analysis, optimization via simulation, nonlinear optimization, and financial and actuarial applications of simulation and optimization methodologies. His e-mail address is ben.feng@uwaterloo.ca. His website is http://www.math.uwaterloo.ca/∼mbfeng/.