Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

EFFICIENT NESTED SIMULATION OF TAIL RISK MEASURES

Jessica O. Dang
Ben M. Feng
Mary R. Hardy

Department of Statistics and Actuarial Science
University of Waterloo
200 University Avenue West
Waterloo, Ontario, N2L 3G1, CANADA

ABSTRACT

Tail risk estimation for portfolios of complex financial instruments is an important enterprise risk management task. Time consuming nested simulations are usually required for such tasks: The outer loop simulates the evolution of risk factors, or the scenarios. Inner simulations are then conducted in each scenario to estimate the corresponding portfolio losses, whose distribution entails the tail risk of interest. In this paper we propose an iterative procedure, called Importance-Allocated Nested Simulation (IANS), for tail risk estimation. We tested IANS in a multiple-period nested simulation setting for an actuarial application. Our numerical results show that IANS can be an order of magnitude more accurate that a standard nested simulation procedure.

1 INTRODUCTION

Tail risk estimation is an important task in financial risk management, as evidenced in its ubiquitous presence in regulations for financial and actuarial industries: Solvency II (EIOPA 2014) applies to most insurers in the European Union, which requires evaluating the 99.5% Value at Risk (VaR) of the change in surplus each year. In North America, insurers use a similar approach under their Own Risk Solvency Assessment (ORSA) obligations (NAIC 2014), where the Conditional Tail Expectation (CTE), also known as the Expected Shortfall (ES), is the tail risk measure of interest.

Under realistic economic models and reasonable variety of financial instruments, tail risk estimation, i.e. estimation of $\rho(L)$ where ρ is the risk measure and $L = g(Y, \mathbf{E}[X|Y])$, usually calls for nested simulation: The outer level simulation projects the underlying risk factors Y, such as values of market indexes and interest rates, to a future time, such as 1 day or 1 year. These outer level sample paths are also known as the *scenarios*. Then inner level simulations are conducted to determine the portfolio values $\mathbf{E}[X|Y]$, which give rises to the portfolio loss L in each scenario based on loss function g. Finally, the empirical distribution of the nested-simulated portfolio losses is then used to estimate the tail risk measure of interest.

In this article we focus on estimating ES efficiently via simulation, although the simulation procedure we introduce here is also applicable to estimation of other tail risk measures such as VaR. ES can be viewed as the average of the largest losses in the tail of the loss distribution. For example, given M equally likely scenarios with losses $L_1, \ldots, L_M$ and a tail probability α (e.g., $\alpha = 0.95$), the $100\alpha\%$-ES is given by

$$ES_\alpha = \frac{1}{(1-\alpha)M} \sum_{j=\alpha M+1}^{M} L_{(j)}. \tag{1}$$

where $L_{(j)}$ is the jth smallest loss, which is also known as the *jth order statistic* of the portfolio losses. We refer to the $(1-\alpha)M$ scenarios whose losses appear in (1) as the *tail scenarios* as they belong to the tail of the loss distribution; the other scenarios are referred to as the *non-tail scenarios*.

Nested simulation is notoriously known for its excessive computational burden, which hinders its applicability in practice, where computational power and run time are constrained. To address this computational challenge, Gordy and Juneja (2010) studied the optimal allocation of a given simulation budget between the outer and inner simulation. Broadie et al. (2011) proposed algorithms to adaptively allocate the given simulation budget to inner simulations. Liu and Staum (2010) used stochastic kriging (Ankenman et al. 2010) to approximate the portfolio loss as a function of the risk factors to improve the estimation efficiency. In addition, the regression method proposed by Broadie et al. (2015) and the kernel smoothing method proposed by Hong et al. (2017) both aim to reduce the number of inner simulations. Readers are encouraged to refer to Hong et al. (2014) for a review on simulation methods in tail risk estimation.

We address the computational challenge of tail risk estimation via nested simulation by first identifying some highly likely tail scenarios (HLTS), on which the available computational budget will be concentrated. The HLTS are identified through an innovative application of simple proxy models and concomitant of order statistics (COS) (David 1973). Suppose $(U_i,V_i), i=1,\ldots,n$ are n independent pairs of r.v. with a common bivariate distribution. We denote the rth order statistic among n U_i's as $U_{r:n}$. Then the V-variate associated with $U_{r:n}$ is called the concomitant of the rth order statistic (David et al. 1977), and is denoted as $V_{[r:n]}$. The rank of $V_{[r:n]}$ among all n V_i's are denoted as $R_{r:n}$. In other words, $R_{r:n}=s$ implies $V_{[r:n]}=V_{s:n}$. The term *concomitant of order statistics* was coined by David (1973) and was independently studied by Bhattacharya (1974) with a different name *induced order statistics*. One well-known application (COS) is to select the best k objects, i.e., V_i's, out of n candidates based on some auxiliary observations, i.e., $\widehat{V}_i$'s (Yeo and David 1984). In our context, we aim to select $(1-\alpha)M$ tail scenarios associated with the true inner simulation model, with high confidence, on the basis of the portfolio losses associated with the proxy model.

Variable annuity (VA) are insurance contracts that are widely used to provide retirement income. Its annual sales in the U.S. market in recent years has been around \$100 billion (LIMRA Secure Retirement Institute 2017). In addition to death and survival benefits to the policyholder, a VA contract also offers benefits that are linked to the performance of certain financial assets, e.g. stock and bonds. As insurance products, such equity-linked benefits often have embedded guarantees to protect the policyholder from downside market risks. From the insurer's perspective, these guarantees can be viewed as embedded financial options. The complexity of the embedded options varies: from European put options of standard Guaranteed Minimum Maturity Benefits (GMMB) to complex combinations of path dependent, exotic lookback and tandem options of the Guaranteed Minimum Income Benefit (GMIB) studied by Marshall et al. (2010). Similar to the risk management of financial options, in practice insurers commonly use a dynamic hedging strategy to mitigate the financial risks of the guarantees in a VA contract, where a hedging portfolio is set up and rebalanced periodically at some rebalancing times. For a comprehensive review of different types of guarantees offered in VA contracts, the hedging of these embedded guarantees, and the modeling of the contracts, we refer readers to Hardy (2003). For ease of exposure, we focus on standard GMMB contracts in this article. The proposed procedure, however, can be applied to more complicated VA contracts.

VAs are often long-term policies, so simple economic models such as log-normal returns with constant volatility are often insufficient for risk management purpose. Moreover, it is a standard practice for insurers to dynamic hedging programs to mitigate the underlying risks of VA contracts. As a result, assessing the tail risks such dynamic hedging program requires a *multi-period nested simulation* procedure, which is considerably more complicated than the *single-period nested simulation* studied in most literature. We consider a multi-period nested simulation setting in our numerical experiments.

The remainder of this article is organized as follows: Section 2 provides an overview of the dynamic hedging practice of VA contracts and the process of a standard nested simulation. Section 3 presents an Important-Allocated Nested Simulation procedure that uses proxy modeling and the rank of concomitant to find the cut-off for potential tail scenarios. Section 4 illustrates the performance of the proposed method in numerical experiments. Section 5 concludes.

2 MOTIVATING EXAMPLE: DYNAMIC HEDGING VIA NESTED SIMULATION

In this section we present the settings, notations, and assumptions through a motivating example: estimation of the ES of the hedging error for a delta-hedged financial instrument (which could itself be an instrument or a portfolio of instruments) whose value is affected by some underlying risk factors. This example is also relevant to VA as many embedded options in a VA guarantee can be modeled by some financial instruments under some simplifying assumptions. In addition, this example illustrates the multi-period nested simulation that we are interested in this article.

In a dynamic hedging program, a hedging portfolio is set up using stocks, bonds, futures and other derivatives to offset or mitigate the risks of the instrument due to fluctuations of the underlying risk factors. For example, a vanilla European put option whose underlying asset is a stock may be hedged by a short position in the underlying stock and a long position in some risk-free bonds. The hedging portfolio is rebalanced periodically, responding to changes in the risk factors. For simplicity, in this article we consider delta hedge programs, in which the composition of the hedging portfolio at time t is determined by the sensitivity of the instrument's value with respect to the risk factor at that time, i.e., the delta at time t, $\Delta(t)$.

Let T be the maturity of the instrument of interest and let $S(t)$ be the value of the relevant risk factors at any time $0 \leq t \leq T$. For simplicity we consider one underlying asset $S(t)$ but the following discussions can easily be generalized to multiple risk factors. We assume that these the underlying asset can be long or short by any amount at any time. Without loss of generality, let $t = 1, 2, \ldots, T$ be the periodic rebalancing times for the hedging portfolio. In a delta hedging program, at time t the hedging portfolio is composed of $\Delta(t)$ units in $S(t)$, and amount $B(t)$ in a risk free zero coupon bond maturing at T. Then, for $t = 0, 1, \ldots, T$, the value of the hedging portfolio at t is given by

$$H(t) = \Delta(t)S(t) + B(t). \tag{2}$$

In a special case $S(t)$ follows a Geometric Brownian Motion, i.e. the Black-Scholes economy, by the risk-neutral pricing theory the price of a vanilla European put option would then be

$$H(t) = Ke^{-r(T-t)}\Phi(-d_2) - S(t)\Phi(-d_1), \tag{3}$$

where K is the strike price, T is the maturity date, r is the per-period continuously compounded risk free rate, and $\Phi(x)$ is the cumulative density function of the standard Normal distribution and

$$d_1(t,T) = \frac{\ln\left(\frac{S(t)}{K}\right) + \left(r + \sigma^2/2\right)(T-t)}{\sigma\sqrt{T-t}} \qquad d_2(t,T) = d_1(t,T) - \sigma\sqrt{T-t}. \tag{4}$$

In addition, the hedging portfolio for the put option in this special case consists of $\Delta(t) = -\Phi(-d_1)$ units of the underlying asset and $B(t) = H(t) - \Delta(t)S(t)$ invested in a zero coupon bond maturing at time T.

Regardless of the model for $S(t)$, suppose the per-period continuously-compounded interest rate is a constant, denoted by r. Then at the end of the tth period (or equivalently at time $t+1$), the value of this hedging portfolio is given by

$$H^{BF}(t+1) = \Delta(t)S(t+1) + B(t)e^r \tag{5}$$

where the superscript "BF" denotes the *hedge brought forward (from time t)* and the argument $(t+1)$ indicates its value at time $t+1$. The hedging error incurred at each rebalancing time t is the difference between the cost of the hedging portfolio at that and the value of the hedging portfolio brought forward from the previous period, i.e.,

$$HE(t) = H(t) - H^{BF}(t), \quad t = 1, \ldots, T. \tag{6}$$

The costs to set up the initial hedging portfolio, $H(0)$, and the present value of the periodic hedging errors, $H(t)$ for $t = 1, \ldots, T$, are recognized as the profit and loss of the hedging program, which is the loss random variable to which we apply a suitable risk measure. Mathematically,

$$L = H(0) + \sum_{t=1}^{T} e^{-rt} HE(t). \tag{7}$$

In general, a multi-period nested simulation is needed to estimate the ES of the hedging loss random variable in (7). A standard multi-period nested simulation process is detailed in the following section.

2.1 Standard Multi-period Nested Simulation Procedure

In a multi-period nested simulation, the outer-level simulation generates M sample paths of risk factors under the real-world measure (or the $\mathbb{P}$-measure), denoted by $S_j := \{S_j(t), t = 0, \ldots, T\}$, for $j = 1, \ldots, M$. Each sample path has T realized values of risk factors at different times. Given the realized values of risk factors at time t in the jth sample path, $S_j(t)$, inner-level simulations are conducted to estimate the corresponding $H_j(t)$, $\Delta_j(t)$, and $B_j(t)$. In each of the N inner simulation at time t, the sample path is initialized with $S_{ij}(t) = S_j(t)$, for $i = 1, \ldots, N$. Then inner simulations are performed to obtain i.i.d samples of $H_{ij}(t)$ and $\Delta_{ij}(t)$. More specifically, the hedging portfolio value $H_{ij}(t)$ can be estimated by Monte Carlo under the risk-neutral measure (or the $\mathbb{Q}$-measure). The sensitivity or Greeks, e.g. $\Delta_{ij}(t)$, can be estimated by the Infinite Perturbation Analysis (IPA) method (Broadie and Glasserman 1996; Glasserman 2013), the likelihood ratio method (L'Ecuyer 1990), simultaneous perturbation (Fu et al. 2016), etc. Then the hedging portfolio value, the delta and the hedge bond value at time t on outer-level sample path j are estimated as

$$H_j(t) = \frac{1}{N}\sum_{i=1}^{N} H_{ij}(t), \qquad \Delta_j(t) = \frac{1}{N}\sum_{i=1}^{N} \Delta_{ij}(t), \qquad B_j(t) = H_j(t) - \Delta_j(t) S_j(t).$$

When inner simulations for all $t = 0, \ldots, T$ are completed for the jth sample path, then the hedging error in the jth scenario, L_j, can be calculated based on (2), (5) – (7). Finally, given a prescribed confidence level α, the ES_α is then estimated by (1) using the simulated hedging losses for all M scenarios.

Compared to the usual single-period nested simulation in other studies, the multi-period nested simulation described above has three nested for-loops: i) for each inner loop simulation i; ii) at each projection time point t; iii) in each outer-level sample paths j. In particular, the inner simulation is performed at *every* time-t along each outer-level sample path, rather than one fixed risk horizon (e.g., 1 day or 1 year). This additional level of simulation further aggravates the computational burden by orders of magnitudes.

The different purposes of the outer- and inner-simulations result in different stochastic asset models being applied. For the purpose of estimating ES of L, a real-world model is used in outer simulations of stock price paths to examine losses associated with the financial instrument under realistic scenarios. Meanwhile, a risk-neutral model is used in inner simulations for estimating $H(t)$, $\Delta(t)$, and $B(t)$ based on risk-neutral pricing theory (Hardy 2003).

3 THE IMPORTANCE-ALLOCATED NESTED SIMULATION (IANS) PROCEDURE

In this section, we present in Algorithm 1 an outline of *Importance-Allocated Nested Simulation (IANS)* procedure for estimating the ES_α of a VA GMMB contract. The IANS procedure replaces the inner simulation steps in the standard nested simulation described in Section 2.1 with a two stage process. Its goal is to achieve higher accuracy in tail risk estimation under a fixed computation budget.

The user must specify some parameters and experiment design choices that govern the behavior of the IANS procedure. For a VA GMMB contract, the put option in a Black-Scholes economy identified in Section 2 is an obvious proxy derivative, as the option payoffs are identical to the guarantee payoffs. To ensure that the proxy losses are highly correlated with the true values under the inner simulation asset model, we dynamically calibrate the Black-Scholes implied volatility in the proxy model to the conditional expected volatility of the real world model, based on the scenario path up to the valuation.

Algorithm 1: Importance-Allocated Nested Simulation of losses for a Delta-hedged VA contract.

Initialization: Simulate M outer scenarios, each is a T-period simulated stock price sample path under the real-world measure.

Stage I: Identification of highly likely tail scenarios (HLTS)

I.1 Select a proxy *financial derivative* and associated *asset model* which provide tractable, analytic hedge costs, and for which the payoff which is expected to be well-correlated to the VA guarantee costs.

I.2 Calibrate the proxy asset model to the underlying risk-neutral asset model in inner-level simulations.

I.3 Implement nested simulation procedure in Section 2.1 but with the analytic hedge calculations for the proxy derivative and asset model replacing the inner simulation step.

I.4 Identify $M^H = (1-\xi)M$ HLTS with the largest simulated proxy loss in Step I.3 for some $\xi \in [0, \alpha]$.

Stage II: Nested simulation with concentrated computation budget

II.1 Allocate remaining computational budget to the M^H HLTS.

II.2 Implement the inner simulation step in Section 2.1 with the original risk neutral asset model and VA payoff, but only for the M^H outer scenarios identified in Step I.4.

II.3 Identify the $M^T = (1-\alpha)M$ largest losses based on the inner simulations.

II.4 Compute output as $ES_\alpha = \frac{1}{M^T}\sum_{j=M^H-M^T+1}^{M^H} L_{(j)}$.

See Proposition A.1 and A.2 in Dang et al. (2018) for details of this calibration. Unlike a standard proxy approach, the highly-likely tail scenarios identified in Stage I of Algorithm 1 do not need to accurately assess the loss values for those scenarios – the proxy step is only for ascertaining a ranking of the loss by scenarios. This means that the IANS procedure is expected to perform well as long as the *rankings of losses* between the proxies and original models are highly correlated, even if the *losses* themselves are not.

The proxies selected in Step I.2 cannot perfectly capture the complexities of the original asset model and VA contract of interest, resulting in potential misclassification of tail scenarios. Therefore we select a proxy confidence level ξ in Step I.4 with some safety margin, so that $\alpha - \xi \geq 0$. We use these $M^H = (1-\xi)M$ HLTS to identify the largest $M^T = (1-\alpha)M$ simulated loss based on the inner simulations, assuming that, with high confidence, the **true tail scenarios** are a subset of the M^H HLTS. This proxy confidence level ξ is an experiment design parameter in IANS. If ξ is very small, the likelihood of capturing the true tail scenarios is high, but at the cost of running the inner simulations on a large number of outer scenarios. Given a fixed budget for the inner simulations, this will generate higher mean square errors in the loss values and ES estimates. On the other hand, if ξ is close to α, the inner simulation budget is focused on fewer scenarios, so those included will have more accurate loss estimates, but some tail scenarios will be wrongly omitted because the proxy loss ranking is not comonotonic with the true loss ranking. Hence there is a trade-off between a high likelihood of including the true tail scenarios ($\xi \to 0$) and high concentration of simulation budget in Stage II ($\xi \to \alpha$).

In this article, we studied two methods to determine the proxy confidence level ξ in Step I.4: By choosing a fixed ξ or by choosing a ξ implied by the distribution of Concomitant of Order Statistics (COS).

3.1 IANS Procedure with Fixed ξ

We first explored the option of choosing a fixed proxy confidence level $\xi = 90\%$, or equivalently a safety margin of $\alpha - \xi = 5\%$. This choice was based on experience with the nested simulation experiments of GMMB contracts. We show in Section 4 that this choice of ξ produces satisfactory results.

3.2 IANS Procedure with ξ Implied by the Distribution of Concomitant of Order Statistics (COS)

Even though the fixed proxy confidence level ξ in Section 3.1 works well in our numerical experiments, there is no statistical guarantee of its effectiveness on other asset and liability models. As such, we developed a different approach to determine ξ based on the statistical properties of Concomitant of Order Statistics.

A nested simulation using the IANS procedure generates a sample with $M^H = (1-\xi)M$ pairs of bivariate output, corresponding to the M^H HLTS identified in the IANS procedure. Each pair of output consists of hedging loss estimated by proxy and by inner simulation, denoted by $(\widehat{L}_i, L_i)$, respectively, for $i = 1,\ldots,M^H$. Given this sample, we try to determine if the $M^T = (1-\alpha)M$ true tail scenarios are a subset of the M^H HLTS. For this purpose, we only need to focus on the quantile of proxy loss and inner simulation loss, rather than the actual value of such losses, which simplifies the problem to one with a bivariate uniform distribution. Because the order statistics of the proxy loss among all scenarios are known, we can study the quantile of inner simulation loss as a concomitant of the quantile of proxy loss. We first derive the first and second moment of the rank of concomitant in a bivariate uniform distribution. We then illustrate how they can be applied in the selection of the proxy confidence level ξ in the IANS procedure.

3.3 Rank of Concomitant of Order Statistics

David et al. (1977) derived a general expression for the expected value of $R_{r:n}$, $E[R_{r:n}]$, for any bivariate distribution of (U,V). As discussed in Section 3.2, quantiles of losses, which are the random variables (U,V) we focus on in this study, have a bivariate uniform distribution. This allows us to use the copula function $C(U,V)$ to denote the distribution function of (U,V). We also denote the density function of the copula as $c(U,V)$. Using the results in (David et al. 1977), we derive Proposition 1.

Proposition 1 Suppose (U,V) has bivariate uniform distribution. In a sample with n pairs of (U,V)'s, the expected value of the rank of concomitant of U's rth order statistic is

$$\begin{aligned}&E[R_{r:n}] \\ &= 1 + n\Big(\int_0^1 \left[\int_0^1 C(u,v)c(u,v)dv\right] f_{U_{r-1:n-1}}(u)du + \int_0^1 \left[\int_0^1 \big(v - C(u,v)\big)c(u,v)dv\right] f_{U_{r:n-1}}(u)du\Big)\end{aligned} \tag{8}$$

where $f_{U_{r:n}}(u)$ is the density function of $U_{r:n}$. More specifically, $f_{U_{r:n}}(u) = \frac{n!}{(r-1)!(n-r)!}u^{r-1}(1-u)^{n-r}$.

Furthermore, in O'Connell (1974), the author derived an expression for $E[R^2_{r:n}]$, where U and V are linearly correlated. Following the same methodology as in O'Connell (1974), we derived Proposition 2.

Proposition 2 Suppose (U,V) has bivariate uniform distribution. In a sample with n pairs of (U,V)'s, the second moment of the rank of concomitant of U's rth order statistic is

$$\begin{aligned}E[R^2_{r:n}] = \quad & 3E[R_{r:n}] - 2 + n(n-1) \times \Big(\int_0^1 \left[\int_0^1 \big(C(u,v)\big)^2 c(u,v)dv\right] f_{U_{r-2:n-2}}(u)du \\ &+ \int_0^1 \left[\int_0^1 \big(v - C(u,v)\big)^2 c(u,v)dv\right] f_{U_{r:n-2}}(u)du \\ &+ 2\int_0^1 \left[\int_0^1 C(u,v)\big(v - C(u,v)\big)c(u,v)dv\right] f_{U_{r-1:n-2}}(u)du\Big)\end{aligned} \tag{9}$$

Given Proposition 1 and 2, the variance of $R_{r:n}$ can be easily derived.

3.4 Application in Importance-Allocated Nested Simulation

Given Algorithm 1, the optimal proxy confidence level ξ in the IANS method is the highest ξ such that the M^T true tail scenarios are a subset of the M^H HLTS. This is equivalent to the highest ξ such that, among all M outer loop scenarios, the lowest rank of proxy losses corresponding the largest M^T inner simulation losses is greater than ξM. Since we could not find this optimal ξ without a full nested simulation of all M outer loop scenarios, we attempt to answer a similar question instead. That is, what is the highest value of ξ such that the rank of the concomitant (i.e. the inner simulation loss) of the $(\xi M+1)^{\text{th}}$ proxy loss is higher than αM? We propose the following iterative process to find this answer.

Given M outer level sample paths and their associated proxy losses $\widehat{L}_i$, $i=1,\ldots M$, we start with an initial proxy confidence level $\xi=\xi_0$ that is slightly lower than α, e.g. $\xi_0=92\%$ for $ES_{95\%}$ estimate, and identify the set of $M^H=(1-\xi_0)M$ HLTS with the largest proxy losses $\widehat{L}_i$. We then implement the inner simulation steps in Section 2.1 only on this set of HLTS and collect M^H pairs of bivariate output $(\widehat{L}_i, L_i)$. We convert each bivariate output to its marginal quantiles among the M^H samples. The converted samples are denoted as (U_i,V_i), for $i=1,\ldots,M^H$. With that, we estimate the empirical copula function and density function of the copula of (U,V). Then we construct a 95% one-sided confidence interval for $R_{(\alpha-\xi_0)M:M^H}$, the rank of concomitant of the $(\alpha-\xi_0)M^{\text{th}}$ order statistics of the M^H proxy losses:

$$\Big(-\infty, \quad E[R_{(\alpha-\xi_0)M:M^H}]+1.645\times\sqrt{\text{Var}[R_{(\alpha-\xi_0)M:M^H}]}\Big) \tag{10}$$

where $E[R_{(\alpha-\xi_0)M:M^H}]$ and $\text{Var}[R_{(\alpha-\xi_0)M:M^H}]$ are calculated by (8) and (9). Let ω denote the upper bound of the confidence interval in (10). If $\omega \geq (\alpha-\xi)M$, then update $\xi=1-\lceil\omega+(1-\alpha)M\rceil \div M$ and redo the above process. The iterative process continues until $\omega<(\alpha-\xi)M$. Then the search for ξ is complete and Stage II of the IANS procedure in Section 3 can be carried out to calculate the ES_α estimate.

The above process searches for the highest ξ in the IANS procedure so that within the M^H pairs of bivariate sample of proxy and inner simulation losses, the upper bound of the confidence interval for the rank of concomitant of one of the best proxy losses is less than $(\alpha-\xi)M$. In other words, we search for the highest ξ such that we could expect with high confidence at least M^T inner simulation losses to be larger than one of the best inner simulation losses in the bivariate sample. We consider the upper bound of the confidence interval for $R_{(\alpha-\xi_0)M:M^H}$ rather than $R_{3:M^H}$, the rank of concomitant of *the best* proxy loss in the sample whose expected value and variance could be derived, in order to reduce the impact from boundary bias when estimating the empirical copula function of the proxy and inner simulation losses.

4 NUMERICAL EXPERIMENTS

To illustrate the performance of the IANS procedure, we use it to estimate $ES_{95\%}$ of VA GMMB hedging loss using different true asset models. We will only present the results from numerical experiments using the Regime-Switching Lognormal asset model as the results using other asset models are similar. Again, we make a few simplifying assumptions so that the payoff of the GMMB contract is identical to that of a European put option on the underlying asset. We consider the risk measure of $ES_{95\%}$, as it is commonly used in valuation and economic capital setting in Canada, consistent with regulatory standards. We assume the GMMB contract has initial fund value and guarantee value of 1,000, 20 years of maturity and monthly rebalancing in the hedging portfolio.

Regime-Switching Lognormal with two regimes is a popular stochastic asset price model for modeling equity-linked insurance contracts (Hardy 2001). The model parameters are provided in Table 1. The financial market is incomplete in the regime-switching model, thus its risk-neutral measure is not unique (Hardy 2001). Given the real-world measure in the regime-switching model, we employ the risk-neutral model studied in Bollen (1998) and Hardy (2001), whose mean conditional log return is $r-\sigma_i^2/2$ for $i=1,2$. All other parameters are the same in the real-world and risk-neutral models.

We assess the IANS procedure using a fixed computational budget for simulation, and compare the accuracy of the resulting ES estimates with those produced with the same simulation budget using the

Table 1: Parameters for the Regime-Switching Model used in Section 4.

(Monthly rate)	Real-World	Risk-Neutral
Risk-free Rate: r	0.002	0.002
Mean - Regime 1 ($\rho = 1$): μ_1	0.0085	0.0013875
Mean - Regime 2 ($\rho = 2$): μ_2	-0.0200	-0.0012000
Standard Deviation - Regime 1: σ_1	0.035	0.035
Standard Deviation - Regime 2: σ_2	0.080	0.080
Transition Probability - from Regime 1: p_{12}	0.04	0.04
Transition Probability - from Regime 2: p_{21}	0.20	0.20

standard nested Monte Carlo (SMC) simulation in Section 2.1. We consider both methods for selecting ξ as described in Section 3.1 and 3.2 when applying the IANS procedure in the numerical experiments.

4.1 Benchmarking Large Scale Nested Simulations

We first conduct a large-scale nested simulation, with 10,000 inner-level simulations and 10,000 outer-level simulations, to obtain an accurate estimate for $ES_{95\%}$. This is a large scale nested simulation because it requires a simulation budget of $10,000 \times 10,000 \times (1 + 12 \times 20) \times (12 \times 20) \div 2 = 2.892 \times 10^{12}$. This estimate serves as a benchmark for assessing the accuracy of other estimators, and is referred to as the *true mean* of $ES_{95\%}$ hereinafter. To illustrate the first stage of the IANS procedure, we replace the inner simulations with closed-form formulas based on the put option proxy derivatives, with the Black-Scholes asset model and examine how many true tail scenarios are correctly identified by the proxies.

Figure 1 depicts the comparisons between the losses that are simulated by standard nested simulation and those by the IANS procedure's proxy simulation. We can see graphically that the values of the inner simulation losses and the proxy losses are highly correlated. This indicates that Stage I in the IANS procedure is able to correctly identify most true tail scenarios without any inner simulation. In fact, 480 out of the 500 true tail scenarios are identified in Stage I in the IANS procedure without any inner simulation. Such robust and accurate identification of tail scenarios leads to the high performance of the IANS procedure, as showcased in subsequent experiments.

4.2 Experiments with Fixed Computation Budget

To demonstrate the efficiency of the IANS procedure using either a fixed ξ or a ξ implied by the concomitant of order statistics (COS), we compare them to three standard nested simulation experiments that use the same simulation budget, but with different allocation between inner- and outer-level simulations, as shown in Table 2. By design, the SMC-5,000-200 experiment has a large number of outer-level simulations and the SMC-200-5,000 experiment has a larger number of inner simulations. The SMC-1,000-1,000 experiment is designed with a more balanced number of inner and outer-level simulation. For the IANS estimators, we set $M = 5,000$. Column (c) in Table 2 indicates the number of outer simulations required. In the experiments using IANS with a fixed $\xi = 90\%$, 500 outer-level simulations were conducted, each with 2,000 inner-level simulations. In the experiments where ξ is implied by the distribution of COS, the number of outer-level simulations varies by experiment, which results in different number of inner-level simulations in each experiment. Nonetheless, the simulation budget consumed in each experiment remains the same. Each experiment design is repeated independently 100 times to produce 100 estimates of $ES_{95\%}$.

Figure 2 depicts the $ES_{95\%}$ estimates in different experiment designs for the GMMB contract. The solid red line in each graph indicates the true mean of $ES_{95\%}$ discussed in Section 4.1. Comparing Figure 2a with Figure 2b and 2c, we see that sufficient number of outer-level simulation reduces the variance while inner-level simulation appears to reduce the bias in estimating tail risk measures. These results are consistent with, for example, Broadie et al. (2011) and Gordy and Juneja (2010). Large number of outer-level

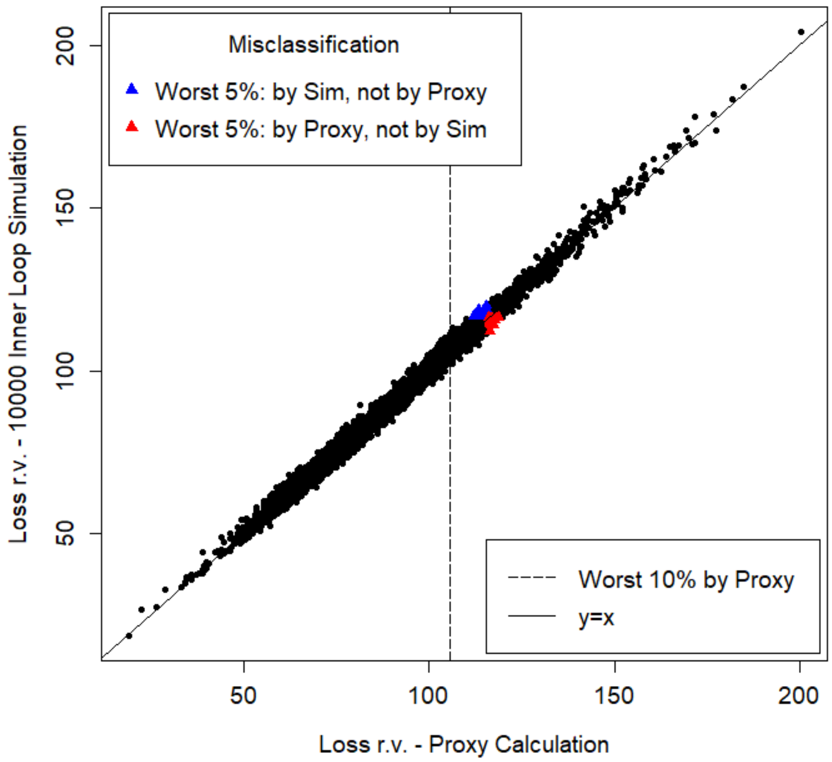

Figure 1: Simulated losses in 10,000 outer scenarios. The x and y coordinates of each point in the figures represent the loss in a scenario, simulated by the IANS proxy simulation and by the true nested simulation, respectively.

simulation reduces the variation in extreme losses simulated from one experiment to another, which reduces the variance of the ES estimate. On the other hand, enough inner-level simulation ensures a more consistent distribution of number of contracts matures in-the-money, which reduces the bias of the simulated loss. Figure 2d and 2e show that the IANS procedure achieves both low bias and low variance compared to the three standard nested simulation experiments using the same simulation budget, regardless of the method used to determine the proxy confidence level ξ.

Table 3 summarizes the relative mean squared errors (RMSEs) for different experiment designs. Each RMSE is calculated as $\frac{1}{100}\sum_{i=1}^{100}\frac{(\widehat{\mu}_i^{est}-\mu)^2}{\mu}$ where $\widehat{\mu}_i^{est}$ is the estimated $ES_{95\%}$ in the ith independent repeated experiment and μ is the true mean of $ES_{95\%}$. The RMSEs are further decomposed into relative bias and relative variance in Table 3 for different experiment designs. Each relative bias is calculated as $\frac{(\frac{1}{100}\sum_{i=1}^{100}\widehat{\mu}_i^{est}-\mu)}{\mu}$, whereas each relative variance is calculated as $\frac{1}{100}\sum_{i=1}^{100}\frac{(\widehat{\mu}_i^{est}-\frac{1}{100}\sum_{i=1}^{100}\widehat{\mu}_i^{est})^2}{\mu}$.

Table 3 demonstrates numerically that, in these examples, the IANS procedure achieves smaller RMSEs compared with straightforward nested simulation using the same simulation budget. The RMSEs indicate the mean squared error in the IANS method experiments are within 2% of the true mean of $ES_{95\%}$ whereas the mean squared error in the SMC experiments are much higher relative to the true mean. Table 3 also shows that both IANS experiments have significantly smaller RMSEs than any of the SMC experiment. In particular, the smaller RMSEs in the IANS experiments are mostly attributed to the smaller relative

Table 2: Simulations in Numerical Experiments.

Experiment Design	(a) M	(b) N	(c) $(1-\xi)M$	Nested Simulation Computational Budget
SMC-5,000-200, $ES_{95\%}$	5,000	200	-	$(a)\times(b)\times\frac{(1+12\times20)\times(12\times20)}{2}=2.892\times10^{10}$
SMC-1000-1000, $ES_{95\%}$	1,000	1,000	-	$(a)\times(b)\times\frac{(1+12\times20)\times(12\times20)}{2}=2.892\times10^{10}$
SMC-200-5,000, $ES_{95\%}$	200	5,000	-	$(a)\times(b)\times\frac{(1+12\times20)\times(12\times20)}{2}=2.892\times10^{10}$
IANS, $ES_{95\%}$ with Fixed $\xi=90\%$	5,000	2,000	500	$(b)\times(c)\times\frac{(1+12\times20)\times(12\times20)}{2}=2.892\times10^{10}$
IANS, $ES_{95\%}$ with COS Implied ξ	5,000	$\frac{2{,}000\times500}{(c)}$	Varies	$(b)\times(c)\times\frac{(1+12\times20)\times(12\times20)}{2}=2.892\times10^{10}$

variance. This echoes the importance of sufficient outer-level simulation relative to inner-level simulations observed in other studies in nested simulations (Broadie et al. 2011; Gordy and Juneja 2010).

Table 3: Relative mean square errors (RMSEs), relative bias and relative variance of ES 95% estimate for different experiment designs.

Experiment Design	RMSE	Relative Bias	Relative Variance
SMC-5,000-200	8.32%	2.34 %	1.20 %
SMC-1,000-1,000	5.38%	0.51 %	5.05 %
SMC-200-5,000	25.50%	-0.24 %	25.43 %
IANS with Fixed $\xi=90\%$	1.28%	0.17 %	1.24 %
IANS with COS Implied $\xi\%$	1.03%	0.03 %	1.03 %

In addition, the IANS experiments using the COS implied proxy confidence level ξ result in a slightly lower RMSE than the IANS experiments using a fixed ξ of 90% because the COS implied ξ ensures a more complete coverage of the true tail scenarios in the IANS procedure. This result highlights the benefit of using the COS implied proxy confidence level ξ to choose the ξ in the IANS procedure: It provides a statistical guarantee that sufficient number of true tail scenarios have been captured in nested simulation. Furthermore, parameters input into the asset and liability models could change frequently in the application of periodic risk reporting, which would give rise to changes in the dynamics of the simulation model. In this case, blindly applying a fixed ξ that worked previously may result in unknowingly missing the true tail scenarios, whereas the COS implied ξ provides a mechanism to avoid such issue.

In addition, we applied the IANS procedure to VaR estimation in the same experiments. We also conducted these experiments on GARCH (1,1) models, as well as more complicated VA contracts, e.g. a Guaranteed Minimum Accumulation Benefit similar to a tandem option (Blazenko et al. 1990) in finance. The results are very similar to those illustrated in this section.

5 CONCLUSION

In this article, we propose an Importance-Allocated Nested Simulation procedure for estimating the ES of loss from VA dynamic hedging strategy. The IANS procedure takes advantage of the special structure of the ES by first identifying a small set of potential tail scenarios based on a proxy for liabilities calculated from a closed-form solution, and then focus the simulation budget on only those scenarios. Based on theories from order statistics, we also propose a prudent and structured method for finding sufficient number of potential tail scenarios to be included in the IANS procedure. We conduct extensive numerical experiments on GMMB contracts. The numerical results show significant improvement in efficiency using the IANS procedure compared to a standard nested simulation.

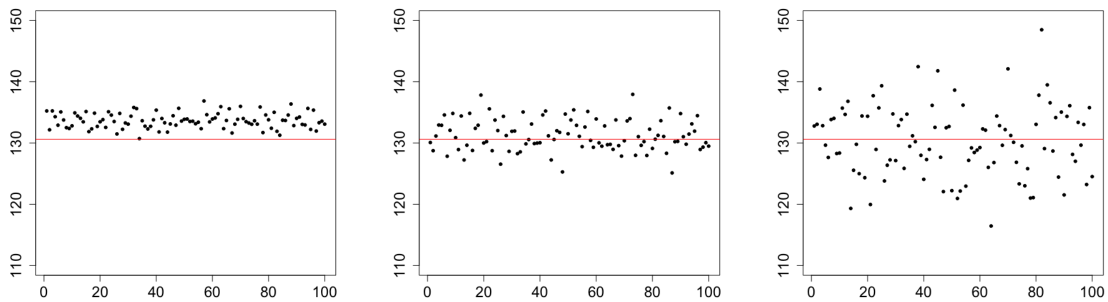

(a) SMC-5,000-200. (b) SMC-1,000-1,000. (c) SMC-200-5,000.

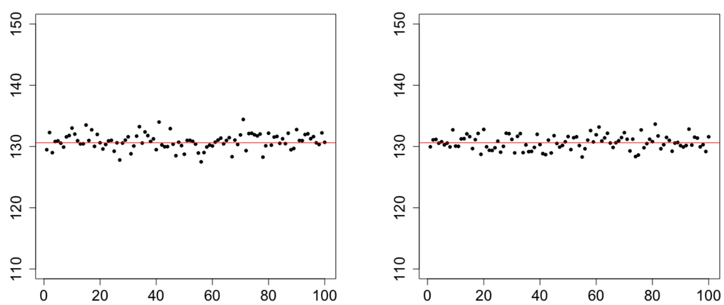

(d) IANS with Fixed ξ. (e) IANS with Concomitant Implied ξ.

Figure 2: Estimated $ES_{95\%}$ of simulated GMMB losses under Regime-Switching Model in 100 independent repeated experiments. The solid red line in each graph indicates the true value estimated by the large scale simulation discussed in Section 4.1.

For future work, we will expand the IANS procedure to more sophisticated VA contract such as GMIB and GMWB, as well as other financial instruments. We will also look for a more general approach to identify the tail scenarios that could accommodate vastly different asset and liability models.

ACKNOWLEDGMENTS

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), funding reference number 203271 (Hardy) and 03755 (Feng). This work was also supported by the Society of Actuaries through a Center of Actuarial Excellence (CAE) Research Grant, and through the Hickman Scholarship held by Jessica Dang. Jessica Dang also receives support from the Ontario Graduate Scholarship.

REFERENCES

Ankenman, B., B. L. Nelson, and J. Staum. 2010. "Stochastic Kriging for Simulation Metamodeling". *Operations Research* 58(2):371–382.

Bhattacharya, P. 1974. "Convergence of Sample Paths of Normalized Sums of Induced Order Statistics". *The Annals of Statistics*:1034–1039.

Blazenko, G. W., P. P. Boyle, and K. E. Newport. 1990. "Valuation of Tandem Options". *Advances in Futures and Options Research* 4:39–49.

Bollen, N. P. 1998. "Valuing Options in Regime-Switching Models". *The Journal of Derivatives* 6(1):38–49.

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.

Broadie, M., and P. Glasserman. 1996. "Estimating Security Price Derivatives Using Simulation". *Management science* 42(2):269–285.

Dang, O., M. Feng, and M. R. Hardy. 2018. "Efficient Nested Simulation for Conditional Tail Expectation of Variable Annuities". *North America Actuarial Journal*. Forthcoming https://ssrn.com/abstract=3045430.

David, H. 1973. "Concomitants of Order Statistics". *Bulletin of the International Statistical Institute* 45(1):295–300.

David, H., M. O'Connell, and S. Yang. 1977. "Distribution and Expected Value of the Rank of a Concomitant of an Order Statistic". *The Annals of Statistics*:216–223.

EIOPA 2014. "The Underlying Assumptions in the Standard Formula for the Solvency Capital Requirement Calculation". Technical report, European Insurance and Occupational Pensions Authority.

Fu, M. C. et al. 2016. *Handbook of Simulation Optimization*. Springer.

Glasserman, P. 2013. *Monte Carlo Methods in Financial Engineering*, Volume 53. Springer Science & Business Media.

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

Hardy, M. 2003. *Investment Guarantees: Modeling and Risk Management for Equity-Linked Life Insurance*, Volume 215. John Wiley & Sons.

Hardy, M. R. 2001. "A Regime-Switching Model of Long-Term Stock Returns". *North American Actuarial Journal* 5(2):41–53.

Hong, L. J., Z. Hu, and G. Liu. 2014. "Monte Carlo Methods for Value-at-Risk and Conditional Value-At-Risk: A Review". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 24(4):22.

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.

L'Ecuyer, P. 1990. "A Unified View of the IPA, SF, and LR Gradient Estimation Techniques". *Management Science* 36(11):1364–1383.

LIMRA Secure Retirement Institute 2017. "Annuity Sales Estimates - 2008-2017".

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

Marshall, C., M. R. Hardy, and D. Saunders. 2010. "Valuation of a Guaranteed Minimum Income Benefit". *North American Actuarial Journal* 14(1):38–58.

NAIC 2014. "NAIC Own Risk and Solvency Assessment (ORSA) Guidance Manual". Technical report, National Association of Insurance Commissioners.

O'Connell, M. J. 1974. "Theory and Applications of Concomitants of Order Statistics".

Yeo, W., and H. David. 1984. "Selection Through an Associated Characteristic, with Applications to the Random Effects Model". *Journal of the American Statistical Association* 79(386):399–405.

AUTHOR BIOGRAPHIES

Jessica O. Dang is a Ph.D. Candidate in Actuarial Science at the University of Waterloo, Canada. Her research interests include efficient simulation algorithms in insurance and finance applications, and design of life insurance and retirement portfolios. She is a Fellow of the Society of Actuaries and a Fellow of the Canadian Institute of Actuaries. Her email address is jessica.dang@uwaterloo.ca.

Ben M. Feng is an Assistant Professor at the University of Waterloo, Canada. He holds a PhD in Industrial Engineering and Management Sciences from Northwestern University, USA. His research interests include quantitative risk management, financial engineering, Monte Carlo simulation design and analysis, and nonlinear optimization. His email address is ben.feng@uwaterloo.ca. His website is http://www.math.uwaterloo.ca/ mbfeng/.

Mary R. Hardy is a Professor at the University of Waterloo, Canada. Her research interests cover aspects of solvency and risk management for both life insurance and pension plans. She is on the editorial board of ASTIN Bulletin and was previously the editor of North American Actuarial Journal and Annals of Actuarial Science. Professor Hardy is a Fellow of the Society of Actuaries (SOA) and of the Institute of Actuaries, and is a Chartered Enterprise Risk Analyst. She was elected to the Board of Governors for the SOA in 2004 and was elected vice-president for 2007-2009. Her email address is mary.hardy@uwaterloo.ca.