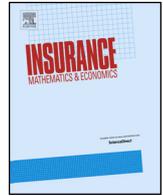




ELSEVIER

Contents lists available at ScienceDirect

# Insurance: Mathematics and Economics

journal homepage: [www.elsevier.com/locate/ime](http://www.elsevier.com/locate/ime)

## Two-stage nested simulation of tail risk measurement: A likelihood ratio approach

Ou Dang<sup>a,\*</sup>, Mingbin Feng<sup>b</sup>, Mary R. Hardy<sup>b</sup><sup>a</sup> Insurance Risk and Finance Research Centre, Nanyang Business School, Nanyang Technological University, Singapore<sup>b</sup> Department of Statistics and Actuarial Science, University of Waterloo, Canada

### ARTICLE INFO

#### Article history:

Received August 2021

Received in revised form August 2022

Accepted 27 October 2022

Available online 3 November 2022

#### JEL classification:

C15

C53

C63

G22

G32

G52

#### Keywords:

Nested simulation

Likelihood ratio method

Importance sampling

Enterprise risk management

Conditional tail expectation

Tail value-at-risk

Expected shortfall

GMWB

### ABSTRACT

Estimating tail risk measures is an important task in many financial and actuarial applications and often requires nested simulations, with the outer simulations representing real world scenarios, and the inner simulations typically used for risk neutral pricing or conditional risk measurement. The standard nested simulation method is highly flexible, able to incorporate complex asset models and exotic payoff structures, but it is also computationally highly burdensome, particularly in a multi-period setting, when inner simulation paths are required at each time step of each outer level scenario. In this study, we propose and analyze a two-stage simulation procedure that efficiently estimates the conditional tail expectation of cost of a dynamic hedging program for a Variable Annuity Guaranteed Minimum Withdrawal Benefit (GMWB), under a multi-period nested simulation. In each of the two stages, the method re-uses the same set of inner level simulation paths for each outer scenario at each time point, using a likelihood ratio method to re-weight the probabilities of each individual path for the different outer scenarios. Our numerical study shows that our two-stage, likelihood ratio weighted procedure can offer a very significant improvement in efficiency, of the order of 95% as measured by the RMSE, compared with a standard nested simulation with the same computational cost.

© 2022 Elsevier B.V. All rights reserved.

### 1. Introduction

Estimating tail risk measures for complex financial derivatives is an important risk management task, often requiring nested simulation. In a standard two-tier simulation, the outer simulation is used to generate paths of the underlying risk factors. These are typically simulated under the real-world probability measure and are known as the *scenarios*. The second tier is the inner simulations, which may be used to estimate the value of the derivative payoff, conditional on the outer scenario path. Typically in financial and actuarial valuation of hedging costs, the inner sample paths are generated under a risk-neutral measure.

Nested simulation is highly flexible and adaptable, and has therefore become an important risk management tool in both financial and actuarial applications, where complex asset models or complex payoff structures make analytic approaches infeasible. However, it can be extremely computationally burdensome, creating a barrier to its application when results are needed at short notice, or where there a large number of model points, for example, within a seriatim valuation of an insurance portfolio.

In this paper, we are concerned with the evaluation of tail risk measures for Variable Annuity (VA) guarantees. VAs are long-term insurance contracts that are widely used for wealth accumulation and for providing retirement income, with annual US sales of around \$100 billion (LIMRA, 2019). A VA contract is very similar to a mutual fund investment, but with additional guarantees and options. The contract premium is invested into a sub-account. The insurer layers additional benefits in the form of guaranteed minimum payouts which protect policyholders from downside market risks. The guarantees are funded through regular deductions from the sub-account. From the

\* Corresponding author.

E-mail address: [ou.dang@ntu.edu.sg](mailto:ou.dang@ntu.edu.sg) (O. Dang).

insurer's perspective, the guarantees can be viewed as embedded financial options, creating significant market risk exposure that can be mitigated using dynamic hedging programs. However, the hedging programs will not be self-financing. Discrete hedging, asset model basis risk, complex guarantee structures and very long terms to maturity all contribute to potentially significant costs arising from hedge rebalancing. Determining tail risk measures for the hedge costs presents a particular computational challenge, as the guarantees are often far out-of-the-money. In this case, a large number of inner simulations are required to accurately evaluate the out-of-the-money options, and a large number of outer scenarios are required to accurately describe the tail of the loss distribution.

Methods to mitigate the computational burden of nested simulation have attracted much research attention. Broadie et al. (2011) propose an algorithm that sequentially allocates a simulation budget to different outer scenarios, one inner simulation at a time. The high level idea of concentrating the inner simulation computation budget on tail scenarios is an inspiration for our work, although we are working with a different problem and a different risk measure from that of Broadie et al. (2011), who use nested simulation to estimate the probability that the loss exceeds a given threshold. Their approach is to allocate the computational budget disproportionately to scenarios just below or above the given threshold, in order to add precision at the key point of the loss distribution. This method is not directly applicable for estimation of tail risk measures such as Value-at-Risk (VaR) or Conditional Tail Expectation (CTE). In the case of VaR estimation, the key threshold is unknown, and indeed is precisely the target of the estimation. In the case of CTE estimation, the entire tail loss distribution above the unknown threshold is critical for accurate estimation of the risk measure.

Other examples from finance include Lee and Glynn (2003), and Gordy and Juneja (2010), who consider the optimal allocation of fixed simulation budgets between the outer and inner simulations by analyzing the asymptotically optimal rates of convergence for different estimators, Broadie et al. (2015), who replace inner simulations by regression estimators, and Hong et al. (2017) who replace inner simulations with a kernel smoothing approach. In actuarial science, recent studies applying efficient nested simulation include Gan and Lin (2015), Lin and Yang (2020), and Feng et al. (2020), which are all concerned with the dual problem of reducing the model points within a large portfolio, as well as estimating the tail risk from that portfolio. Our work takes a different approach to each of these. We treat the outer scenarios as fixed, and focus the computational budget allocation on the inner simulation step. This has some advantage in a life insurance context, where the outer scenarios may be centrally generated and used for multiple portfolios. We do not eliminate or replace the inner simulation step; instead, we propose a method for increasing the efficiency by reusing inner simulations across multiple scenarios, and we focus on methods for accurate and efficient valuation of the tail risk measure of a single contract. Our approach may be combined with cluster analysis for selecting representative contracts to address the dual problem.

Given a set of outer scenarios, tail risk measure estimation can be viewed as a two stage process. The first stage is to identify the tail scenarios from the full set of outer scenarios, and the second stage is to accurately and efficiently estimate the losses in the tail scenarios. Other work using this two-stage approach, with a given set of outer scenarios, includes Lan et al. (2010), who propose an iterative screening procedure to eliminate scenarios that are unlikely to have large losses, Liu and Staum (2010), who use stochastic kriging to select the tail scenarios from the full set, and Dang et al. (2020) and Dang et al. (2022), who identify likely tail scenarios based on a proxy analytic evaluation, saving most of the computation budget for the second stage. Each of these approaches is limited in its applicability; for example, the Lan et al. (2010) method requires the standard deviation of the losses to be relatively small, and Dang et al. (2022) require the existence of an effective proxy valuation model, which will not be available for some of the more complex dynamic guarantees embedded in VA policies. In this paper our approach is very general, which means that it can be applied more widely than previous methods. For both stages, we use simulation. In the first stage, we use a small part of the simulation budget to achieve a rough valuation, sufficient to identify the likely tail scenarios. In the second stage, the remaining simulation budget is applied to the scenarios that are most likely to generate tail losses, to ensure that the tail risk measure is accurately evaluated for the given set of outer scenarios. The inner simulation efficiency in both stages is improved by using the likelihood ratio estimator, which allows the simulated inner sample paths to be re-used across different outer scenarios. Reusing simulation outputs to improve efficiency is the key idea of green simulation, which was first proposed by Feng and Staum (2017) and further developed by Feng and Staum (2021). Our work builds on theirs, but with a very different focus. Their work considered repeated experiments whose parameters are driven by an ergodic stochastic process, with an analysis of the asymptotic convergence of the estimator. Our focus is more applied; in particular, our aim is to increase efficiency in the context of inner simulations which are parametrized by the given outer scenarios.

The likelihood ratio (LR) method (also known as the score function method), was first studied by Beckman and McKay (1987), and has subsequently been used in a wide range of applications, including metamodeling, sensitivity analysis, and optimization. See, for example, Kleijnen and Rubinstein (1996) and Maggiar et al. (2018), who use LR for optimization of complex computational simulation results, and Glasserman and Xu (2014), who use LR to examine and quantify model risk in the context of portfolio management. Concurrently with our work, Feng and Li (2022) developed an application of the LR method to variable annuities. Their method differs from ours in requiring the user to set blocks of scenarios, such that inner simulations are only shared within blocks, to ensure that variances are limited. Our method manages the problem of unlimited variance using a mixture likelihood approach, described in Section 3, which eliminates the need for the user to identify scenario blocks.

The VA problem that we consider requires a multi-period nested simulation, as the inner simulation step is repeated at each hedge rebalancing date. This means that many of the single-period methods developed in the finance literature cannot be used. Fig. 1 illustrates the difference between the single period problem and the multi-period problem. If a 30-year contract is dynamically hedged monthly, then nested simulation requires  $T = 360$  inner simulations for each scenario, compared to a single inner simulation required for the single period problem.

This study has two contributions:

1. We develop a two-stage procedure for tail risk estimation that uses the likelihood ratio method to reuse inner simulation paths, significantly improving efficiency compared with a standard nested simulation. The first stage uses small-scale inner simulations to identify a set of highly likely tail scenarios. In the second stage, the remaining computational budget is concentrated on the scenarios identified in the first stage, again using likelihood ratio estimators to pool inner samples from different scenarios to improve efficiency.
2. We adapt the likelihood ratio estimators to the special case of the Guaranteed Minimum Withdrawal Benefit (GMWB) (described in Section 4), which creates an unusual challenge arising from (i) the very complex nature of the guarantee, and (ii) the possibility that the policyholder's fund becomes fully depleted.

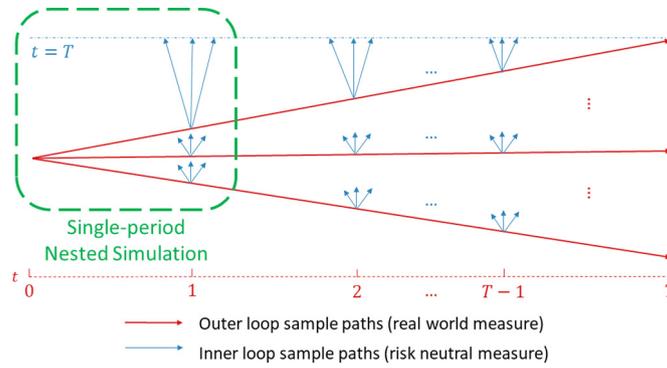


Fig. 1. Nested simulation structure. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

3. We briefly discuss how the likelihood ratio estimators can be applied to a portfolio of VA contracts, with a mixture of GMWB and Guaranteed Minimum Maturity Benefit (GMMB).

The rest of this paper is organized as follows. In Section 2 we present the mathematical settings for the multi-period nested simulation problems studied in this paper. In Section 3 we present details of the two-stage simulation procedure using likelihood ratio estimators. In Section 4 we adapt the two-stage procedure to estimate the tail risks of the dynamically hedged GMWB. In Section 5 we examine the performance of the two-stage procedure for the GMWB numerically. In Section 6 we illustrate with a numerical example how the two-stage procedure is applied to the nested simulation of a portfolio of GMWB and GMMB contracts. Section 7 concludes this paper and summarizes a few possible future lines of research.

## 2. Mathematical settings and problem statement

Consider a variable annuity contract whose embedded option value depends on some state variables, such as equity returns, interest rates, demographics or loss indices. We are interested in estimating a tail risk measure of a loss random variable that depends on the current and future values of this option.

Let  $t=0$  be the current time and let  $T > 0$  be the maturity date of the embedded option. We denote the state variable at time  $t$ , which may be a vector or a scalar, by  $S_t$ , and we assume that the state process  $\{S_t, 0 \leq t \leq T\}$  is modeled by a sufficiently regular stochastic process, defined on a probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ , with a natural filtration  $\mathcal{F}_t$  governing its evolution, where  $\mathbb{P}$  denotes the real-world physical measure. We assume the existence of a risk-neutral measure  $\mathbb{Q}$ , equivalent to  $\mathbb{P}$ , such that financial assets can be valued as the expected discounted payoff under  $\mathbb{Q}$ . We assume the state process is simulated at discrete times  $t = 0, 1, \dots, T$ .

For any  $t = 0, \dots, T$ , we denote the real-world state variables up to time  $t$  by  $\mathbf{S}_t = \{S_0, \dots, S_t\}$ . The entire real-world path from time 0 to  $T$ , i.e.,  $\mathbf{S}_T$ , is the outer scenario.

At each time  $t = 0, \dots, T - 1$  of an outer scenario, an inner simulation is run for the period from  $t + 1$  to  $T$ , under the risk neutral measure, conditional on the outer scenario up time  $t$ , i.e. on  $\mathbf{S}_t$ . The inner sample path process from time  $t + 1$  to time  $T$  is denoted by  $\tilde{\mathbf{S}}_{t+} = \{\tilde{S}_{t+1}, \dots, \tilde{S}_T\}$ .

Given a time- $t$  scenario  $\mathbf{S}_t$  for  $t = 0, \dots, T - 1$ , the inner simulation is performed to estimate the conditional expected output:

$$\Delta_t(\mathbf{S}_t) = \mathbb{E}_t [H_t(\tilde{\mathbf{S}}_{t+}) | \mathbf{S}_t], \tag{1}$$

where  $H_t$  is a time- $t$  loss function that depends on the random path  $\tilde{\mathbf{S}}_{t+}$ , and on the particular contract being valued. At  $T$ , given the outer scenario  $\mathbf{S}_T$  there is no further uncertainty, and

$$\Delta_T(\mathbf{S}_T) = H_T(\mathbf{S}_T)$$

Examples of  $H_t(\tilde{\mathbf{S}}_{t+})$  and  $\Delta_t(\mathbf{S}_t)$  in discrete time hedging applications are provided in Section 2.1.

We assume that the loss random variable of interest can be expressed as a function of all the outputs from  $t = 0, \dots, T$ , and so can be written as

$$L(\mathbf{S}_T) = L(\Delta_0(\mathbf{S}_0), \Delta_1(\mathbf{S}_1), \dots, \Delta_T(\mathbf{S}_T)). \tag{2}$$

In our context, the loss comprises the total discounted hedge costs, net of fee income. The initial cost is the value of the initial hedge portfolio. The hedge cost at each subsequent rebalancing date  $t = 1, 2, \dots, T$  is the difference between the value of time  $t$  hedge portfolio and the time  $(t-1)$  hedge portfolio. For  $t = 1, 2, \dots, T$ , this difference depends only on  $\Delta_t(\mathbf{S}_t)$  and  $\Delta_{t-1}(\mathbf{S}_{t-1})$ . Note that the loss  $L(\mathbf{S}_T)$  may be negative when the fee income exceeds the hedge costs.

As noted above, in this paper we consider a fixed set of outer scenarios. Our concern is on the accuracy of tail risk assessment given the scenario set. Thus, the ‘true’ loss random variable that we are concerned with is that associated with the given set of outer scenarios, and the error measures that we use are relative to the most accurate possible assessment of loss for the given outer scenario set. The overall error will depend on the inner simulation accuracy of the tail scenarios, which is the focus of this paper, and the sampling error from the outer simulation step. Our approach is consistent with how nested simulations are carried out in practice in the life insurance industry. The outer scenarios are typically simulated according to regulatory requirements or internal risk management practices, and may be used for several different portfolios, involving different contract types and durations.

We assume that we have  $M$  given outer scenarios  $\mathbf{S}_{T,i}$ ,  $i = 1, \dots, M$ . We denote the true losses for the  $i$ -th scenario by  $L_i = L(\mathbf{S}_{T,i})$ , and we let  $L_{(1)} < L_{(2)} < \dots < L_{(M)}$  be the corresponding ordered losses (assume no ties for convenience). We are interested in estimating the  $\alpha$ -Conditional Tail Expectation (CTE) (Wirch and Hardy, 1999), which, if  $\alpha M$  is an integer, is given by

$$CTE_\alpha = \frac{1}{(1-\alpha)M} \sum_{i=\alpha M+1}^M L_{(i)} = \frac{1}{(1-\alpha)M} \sum_{i:\mathbf{S}_{T,i} \in \mathcal{T}_\alpha} L_i, \tag{3}$$

where  $\mathcal{T}_\alpha = \{\mathbf{S}_{T,i} : L_i > L_{(\alpha M)}\}$  is the set of scenarios that are included in the calculation of  $CTE_\alpha$  in (3); we call  $\mathcal{T}_\alpha$  the *true tail scenario set*.<sup>1</sup>

In a standard nested simulation, for each time step  $t$  of each scenario, we simulate  $N$  independent and identically distributed inner sample paths to estimate  $\Delta_t(\mathbf{S}_{t,i})$  via the following Monte Carlo estimator:

$$\widehat{\Delta}_t^{NS}(\mathbf{S}_{t,i}) = \frac{1}{N} \sum_{j=1}^N H_t(\widetilde{\mathbf{S}}_{t+,ij}), \quad \widetilde{\mathbf{S}}_{t+,ij} \stackrel{i.i.d.}{\sim} f(\widetilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,i}), \quad \forall j = 1, \dots, N, \tag{4}$$

where  $f(\widetilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,i})$  is the conditional probability density function (pdf) of the inner sample paths, given the outer scenario  $\mathbf{S}_{t,i}$ . We refer to  $\widetilde{\mathbf{S}}_{t+,ij}$  as the  $(ij)$ -th *inner sample path* and  $H_t(\widetilde{\mathbf{S}}_{t+,ij})$  the  $(ij)$ -th *inner simulation output*. A standard multi-period nested simulation is computationally burdensome as it requires  $M \times N \times T$  inner simulation outputs. In some applications, each inner simulation output requires significant computations – in Section 4, we consider discrete time hedging of a complex variable annuity contract, for which the inner simulation outputs are calculated via recursion.

Using  $\widehat{\Delta}_t^{NS}(\mathbf{S}_{t,i})$  in (4) for all  $t = 0, \dots, T$ , the estimated value of loss for scenario  $i$  is

$$\widehat{L}_i^{NS} = L(\widehat{\Delta}_0^{NS}(\mathbf{S}_{0,i}), \widehat{\Delta}_1^{NS}(\mathbf{S}_{1,i}), \dots, \widehat{\Delta}_T^{NS}(\mathbf{S}_{T,i})), \quad i = 1, \dots, M.$$

Denote the ordered estimated losses by  $\widehat{L}_{(1)}^{NS} < \widehat{L}_{(2)}^{NS} < \dots < \widehat{L}_{(M)}^{NS}$ , then the  $CTE_\alpha$  can be estimated by

$$\widehat{CTE}_\alpha^{NS} = \frac{1}{(1-\alpha)M} \sum_{i=\alpha M+1}^M \widehat{L}_{(i)}^{NS} = \frac{1}{(1-\alpha)M} \sum_{i:\mathbf{S}_{T,i} \in \widehat{\mathcal{T}}_\alpha^{NS}} \widehat{L}_i^{NS}, \tag{5}$$

where  $\widehat{\mathcal{T}}_\alpha^{NS} = \{\mathbf{S}_{T,i} : \widehat{L}_i^{NS} > \widehat{L}_{(\alpha M)}^{NS}\}$  is the set of scenarios associated with the  $(1-\alpha)M$  largest simulated losses, and that are therefore included in the calculation of  $\widehat{CTE}_\alpha^{NS}$  in (5).  $\widehat{\mathcal{T}}_\alpha^{NS}$  is the *nested simulation tail scenario set*.

We observe some drawbacks of the standard nested simulation procedure.

- The delta estimate for scenario  $i$ , from equation (4), is an average of only the simulation outputs in scenario  $i$ . The simulation outputs in the other scenarios are ignored, which may be an inefficient use of the computation involved. In Section 3.1, we show how the likelihood ratio method is employed to reuse all available simulation outputs.
- As we are treating the  $M$  outer scenarios as fixed, the difference between the CTE estimates in equation (5) and equation (3) is due solely to the inner simulation sampling error. Specifically, the inner simulation noise affects the accuracy of equation (5) in two ways:
  - (1) Classification of tail scenarios. Due to the inner simulation sampling variability, the estimated losses  $\widehat{L}_i^{NS}$  will differ from true losses  $L_i$ . As a result, the corresponding tail scenario sets may be different, i.e.,  $\widehat{\mathcal{T}}_\alpha^{NS} \neq \mathcal{T}_\alpha$ .
  - (2) Estimation of tail losses for scenarios in the true tail scenario set  $\mathcal{T}_\alpha$ . We see from equation (3) that losses from non-tail scenarios are irrelevant for estimating the CTE.

Our two-stage procedure is specifically designed to identify tail scenarios and accurately estimate tail losses. Very little computation is spent on non-tail scenarios.

### 2.1. Hedging loss in discrete time delta hedging

In this section we describe the simulation model function  $H_t(\widetilde{\mathbf{S}}_{t+})$ , the expected value of the inner simulation output  $\Delta_t(\mathbf{S}_t)$ , and the loss random variable  $L(\mathbf{S}_T)$ , in the context of discrete time hedging for the embedded option of a variable annuity contract.

Let  $v_t(\mathbf{S}_{t+})$  be the value at  $t$  of the future liability beyond time  $t$ , based on the sample path  $\mathbf{S}_{t+}$ . The risk-neutral value of the liability at  $t$ , conditioning on  $\mathbf{S}_t$ , is denoted by  $V_t(\mathbf{S}_t) = \mathbb{E}[v_t(\widetilde{\mathbf{S}}_{t+}) | \mathbf{S}_t]$ . Let  $v_0(\mathbf{S}_{0+}) = v_0(\mathbf{S}_T)$  denote the present value at the contract inception of the embedded option payoff at maturity, for scenario  $\mathbf{S}_T$ . We assume that the insurer uses a delta hedge to mitigate the potential loss. This requires the insurer to hold a hedge portfolio at  $t$  of  $\Delta_t S_t$  in stocks, and  $B_t = V_t(\mathbf{S}_t) - \Delta_t S_t$  in zero coupon bonds, where

$$\Delta_t = \frac{\partial V_t(\mathbf{S}_t)}{\partial \mathbf{S}_t}$$

In perfect conditions we could set up and dynamically rebalance a perfect self-financing hedge, at a cost of  $P_0 = \Delta_0 S_0 + B_0$ , which would perfectly replicate the payoff. In practice, there are many reasons why the hedge portfolio does not provide perfect replication, with the main factors being that the hedge is rebalanced at discrete time intervals, and that there is a basis risk from the pricing model.

<sup>1</sup> Typically, in practice,  $\alpha M$  will be an integer. If it is not, then we can use the floor function  $\lfloor \alpha M \rfloor$ .

In a discrete rebalancing plan, cash infusions or withdrawals arise at each rebalancing date. The hedging error at  $t$  is the difference between the cost of the hedge required at  $t$ , and the value of the hedge brought forward from the previous rebalancing date. The hedging error can be positive or negative. The loss random variable is the initial cost of the hedge portfolio plus the discounted expected value of the hedging errors.

At each time  $t = 0, 1, \dots, T - 1$ , the delta hedge portfolio value is given by

$$P_t = \Delta_t S_t + B_t. \tag{6}$$

This hedge is held for one period and then rebalanced at  $t + 1$ . If we assume a constant interest rate  $r$  in each period, then for  $t = 1, 2, \dots, T$ , the value at  $t$  of the hedge portfolio brought forward from  $t - 1$  is

$$P_t^{bf} = \Delta_{t-1} S_t + B_{t-1} e^r. \tag{7}$$

The hedge error at  $t$  is the difference between the cost of the hedge portfolio, and the value of the hedge brought forward from the previous period, i.e.,

$$HE_t = P_t - P_t^{bf}, \quad t = 1, 2, \dots, T - 1. \tag{8}$$

So we have

$$\begin{aligned} L(\mathbf{S}_T) &= P_0 + \sum_{t=1}^{T-1} e^{-rt} HE_t + [v_0(\mathbf{S}_T) - e^{-rT} P_T^{bf}] \\ &= \sum_{t=0}^{T-1} \Delta_t [e^{-rt} S_t - e^{-r(t+1)} S_{t+1}] + v_0(\mathbf{S}_T). \end{aligned} \tag{9}$$

For complex embedded options, inner simulations are required to estimate the deltas. A common estimation method is the pathwise estimate, or the infinitesimal perturbation analysis (IPA), which is based on the following identity that holds under some conditions (Glasserman, 2013):

$$\frac{\partial V_t(\mathbf{S}_t)}{\partial \mathbf{S}_t} = \frac{\partial}{\partial \mathbf{S}_t} \mathbb{E}[v_t(\tilde{\mathbf{S}}_{t+}) | \mathbf{S}_t] = \mathbb{E} \left[ \frac{\partial}{\partial \mathbf{S}_t} v_t(\tilde{\mathbf{S}}_{t+}) \middle| \mathbf{S}_t \right].$$

Denote  $H_t(\tilde{\mathbf{S}}_{t+}) = \frac{\partial}{\partial \mathbf{S}_t} v_t(\tilde{\mathbf{S}}_{t+})$ , then  $\Delta_t(S_t) = \mathbb{E} [H_t(\tilde{\mathbf{S}}_{t+}) | \mathbf{S}_t]$  so the delta can be estimated via inner simulation. We show the derivation of  $H_t(\tilde{\mathbf{S}}_{t+})$  for a GMWB in Section 4.1.

Although we only consider delta hedging in this study, our procedure is applicable to discrete time hedging programs involving multiple option Greeks, if they can be estimated via inner simulation using, for example, the pathwise estimate.

### 3. Two-stage nested simulation with likelihood ratio estimators

#### 3.1. Reusing simulation outputs using mixture likelihood ratio estimators

The green simulation design paradigm of Feng and Staum (2017) reuses simulation outputs in temporally repeated experiments, so that simulation outputs from previous studies can be used to improve the efficiency in future studies, using likelihood ratio based estimators. The principle can be adapted to the context of nested simulations, which can be viewed as artificial temporally repeated experiments.

Assumption 3.1 ensures that the likelihood ratio estimators in this paper are well-defined and can be calculated.

**Assumption 3.1.** The conditional probability density functions,  $f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,i})$ , are well-defined and can be calculated for all  $i = 1, \dots, M$  and  $t = 0, 1, \dots, T - 1$ . Moreover, for each  $t = 0, 1, \dots, T - 1$ , the supports of  $f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,i})$  for all  $i = 1, \dots, M$  are identical.

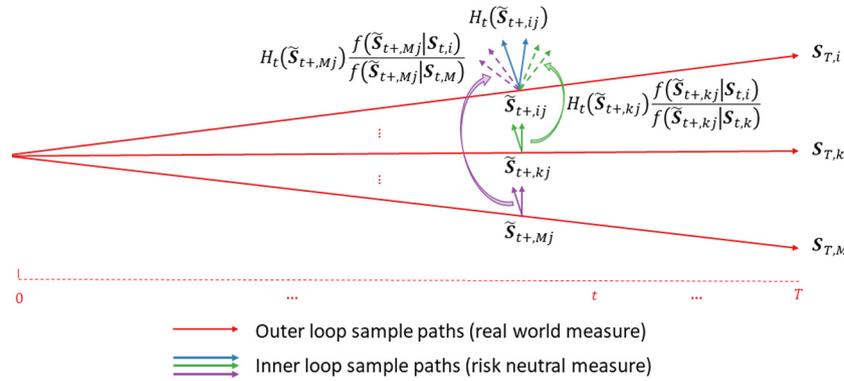
If Assumption 3.1 holds, then we have the following importance sampling identity.

$$\begin{aligned} \Delta_t(\mathbf{S}_{t,i}) &= \mathbb{E} [H_t(\tilde{\mathbf{S}}_{t+}) | \mathbf{S}_{t,i}] \\ &= \int H_t(\tilde{\mathbf{S}}_{t+}) f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,i}) d\tilde{\mathbf{S}}_{t+} = \int H_t(\tilde{\mathbf{S}}_{t+}) \frac{f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,i})}{f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,k})} f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,k}) d\tilde{\mathbf{S}}_{t+} \\ &= \mathbb{E} \left[ H_t(\tilde{\mathbf{S}}_{t+}) \frac{f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,i})}{f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,k})} \middle| \mathbf{S}_{t,k} \right]. \end{aligned} \tag{10}$$

Equation (10) allows us to use the inner simulation outputs from the  $k$ -th scenario,  $\mathbf{S}_{t,k}$ , to estimate the delta output for the  $i$ th scenario,  $\Delta_t(\mathbf{S}_{t,i})$ , using the likelihood ratio (LR) estimator:

$$\hat{\Delta}_{t,k}^{LR}(\mathbf{S}_{t,i}) = \frac{1}{N} \sum_{j=1}^N H_t(\tilde{\mathbf{S}}_{t+,kj}) \frac{f(\tilde{\mathbf{S}}_{t+,kj} | \mathbf{S}_{t,i})}{f(\tilde{\mathbf{S}}_{t+,kj} | \mathbf{S}_{t,k})}, \quad \tilde{\mathbf{S}}_{t+,kj} \stackrel{i.i.d.}{\sim} f(\tilde{\mathbf{S}}_{t+} | \mathbf{S}_{t,k}), \quad \forall j = 1, \dots, N. \tag{11}$$

The intuition behind the LR estimator in equation (11) is that reweighting the  $k$ -scenario's inner simulation outputs  $H_t(\mathbf{S}_{t+,kj})$ ,  $j = 1, \dots, N$  using the likelihood ratios effectively adjusts the probability such that the scenario  $k$  output can be applied to the scenario  $i$



**Fig. 2.** Schematic illustration of reusing simulation outputs via likelihood ratio estimators. Scenario  $S_{T,i}$  represents the target scenario whereas scenario  $S_{T,k}$  and  $S_{T,M}$  represent two different sampling scenarios.

estimate. We refer to the  $k$ -th scenario here as the *sampling scenario* and the  $i$ -th scenario as the *target scenario*. Since we can do this for any  $k$  and  $i \neq k$ , all of the inner simulation paths can be used with each of the outer scenarios, at each time step. This pooling of simulation outputs is expected to greatly improve the estimation accuracy for each target scenario compared to averaging only the simulation outputs in that scenario.

Equation (11) shows that the likelihood ratio method is mathematically identical to importance sampling, but it differs in means and goals. Importance sampling seeks the best sampling distribution for the goal of reducing variance. The likelihood ratio method, however, has no control over the sampling distribution. The goal is to reduce computation; variance reduction is expected but not guaranteed.

Fig. 2 depicts the reuse of simulation outputs: The left to right solid lines represent outer scenarios, while the upward solid lines represent the specific inner sample paths generated for each scenario. The broken lines illustrate how the inner simulations from scenario  $k$  and  $M$  are reused for scenario  $i$ , along with the paths generated specifically for scenario  $i$ .<sup>2</sup> The phrase “reuse simulation outputs” includes two aspects:

1. Reuse the sampling scenario’s simulation outputs,  $H_t(\tilde{S}_{t+,kj})$ ,  $j = 1, \dots, N$ , to estimate the target scenario’s expected output  $\Delta_t(S_{t,i})$ . These outputs are only calculated once but are reused  $M$  times for the  $M$  scenarios. As shown in Fig. 2, reusing simulation outputs effectively increases the number of outputs in the target scenario and is expected to improve accuracy. Note that reusing the simulation outputs does not require any additional inner simulation; likelihood ratio calculations are needed but these generally require less computation than new simulations. Section 3.1.1 elaborates on the computational aspects of the likelihood ratio method.
2. Concatenate the sampling scenario’s inner sample paths  $\tilde{S}_{t+,kj}$  to the target scenario  $S_{t,i}$ , as shown by the upward dashed lines in Fig. 2. Usually, this concatenation is a useful visual aid for the likelihood ratio calculation  $\frac{f(\tilde{S}_{t+,kj}|S_{t,i})}{f(\tilde{S}_{t+,kj}|S_{t,k})}$ . For complex embedded options, such as the GMWB, direct concatenation may be invalid and adaptations are needed before applying the likelihood ratio method. This is discussed further in Section 4.2.

Note that, similarly to importance sampling, the LR estimator can have large or even infinite variance due to the skewness of the likelihood ratio. Specifically, even though the likelihood ratio is positive and has unit expectation, that is, even though  $\frac{f(\tilde{S}_{t+}|S_{t,i})}{f(\tilde{S}_{t+}|S_{t,k})} > 0$

and  $\mathbb{E} \left[ \frac{f(\tilde{S}_{t+}|S_{t,i})}{f(\tilde{S}_{t+}|S_{t,k})} \middle| S_{t,k} \right] = \mathbb{E} [1 | S_{t,i}] = 1$ , the likelihood ratio itself can be unbounded. For example, consider the exponential distribution  $f(x|\lambda) = \lambda e^{-\lambda x}$ . For  $\lambda_2 > \lambda_1 > 0$ , the likelihood ratio  $\frac{f(x|\lambda_1)}{f(x|\lambda_2)} = \frac{\lambda_1}{\lambda_2} e^{-(\lambda_1 - \lambda_2)x} \rightarrow \infty$  as  $x \rightarrow \infty$ . Such a highly-skewed likelihood ratio can lead to LR estimators with infinite variance. In the context of nested simulation, this problem can arise when the sampling scenario is significantly different from the target scenario.

The LR estimator  $\hat{\Delta}_{t,k}^{LR}(S_{t,i})$  in equation (11) uses inner simulation outputs from a single sampling scenario, the  $k$ -th scenario, to estimate the  $i$ -th scenario delta,  $\Delta_t(S_{t,i})$ . Given  $M$  sampling scenarios,  $\hat{\Delta}_{t,k}^{LR}(S_{t,i})$ ,  $k = 1, \dots, M$  are all unbiased estimators of  $\Delta_t(S_{t,i})$ . One way to combine these is to average them, i.e.,

$$\hat{\Delta}_t^{ILR}(S_{t,i}) = \frac{1}{M} \sum_{k=1}^M \hat{\Delta}_{t,k}^{LR}(S_{t,i}).$$

This is the *individual likelihood ratio (ILR)* estimator in Feng and Staum (2017). The ILR estimator’s variance is a weighted sum of the variances of the individual LR estimators  $\hat{\Delta}_{t,k}^{LR}(S_{t,i})$  for  $k = 1, \dots, M$ , therefore the ILR estimator will have infinite variance when one or more of the LR estimators has infinite variance.

A remedy proposed by Feng and Staum (2017) is the *mixture likelihood ratio (MLR)* estimator, which is given by

$$\hat{\Delta}_t^{MLR}(S_{t,i}) = \frac{1}{MN} \sum_{k=1}^M \sum_{j=1}^N H_t(\tilde{S}_{t+,kj}) \frac{f(\tilde{S}_{t+,kj}|S_{t,i})}{\bar{f}_M(\tilde{S}_{t+,kj})}, \tag{12}$$

<sup>2</sup> The individual likelihood ratios are shown in Fig. 2 for illustration only. We use the mixture likelihood ratios in our numerical experiments.

where the *mixture pdf* for the inner paths is defined as

$$\bar{f}_M(\tilde{\mathbf{s}}_{t+}) = \frac{1}{M} \sum_{i=1}^M f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i}). \tag{13}$$

The intuition behind the MLR estimator (12) is similar to that for the LR estimator (11), as both reweight the simulation outputs by the appropriate likelihood ratios, but they differ in the interpretation of the sampling distributions for inner paths. The LR estimator in equation (11) treats the inner sample paths  $\{\tilde{\mathbf{S}}_{t+,kj}, j = 1, \dots, N\}$ , for  $k = 1, 2, \dots, M$  as  $M$  different samples, each following the conditional distribution  $f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,k})$ ,  $k = 1, 2, \dots, M$ . The MLR estimator, in contrast, views  $\{\tilde{\mathbf{S}}_{t+,kj}, k = 1, \dots, M, j = 1, \dots, N\}$  as one sample from the mixture distribution  $\bar{f}_M(\tilde{\mathbf{s}}_{t+})$ .

The MLR estimator has been studied by Veach and Guibas (1995), Feng and Staum (2017), and Elvira et al. (2019). We show here that it is unbiased (like the LR estimator) and that the likelihood ratio for the MLR is bounded (unlike the LR estimator). To show that it is unbiased, recall that the sampling distribution for  $\tilde{\mathbf{S}}_{t+,kj}$  is  $f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,k})$  for all  $j = 1, \dots, N$ , so we have

$$\begin{aligned} \mathbb{E} \left[ \widehat{\Delta}_t^{\text{MLR}}(\mathbf{S}_{t,i}) \right] &= \frac{1}{MN} \sum_{k=1}^M \sum_{j=1}^N \mathbb{E} \left[ H_t(\tilde{\mathbf{S}}_{t+,kj}) \frac{f(\tilde{\mathbf{S}}_{t+,kj} | \mathbf{S}_{t,i})}{\bar{f}_M(\tilde{\mathbf{S}}_{t+,kj})} \right], \quad \tilde{\mathbf{S}}_{t+,kj} \sim f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,k}) \\ &= \frac{1}{MN} \sum_{k=1}^M \sum_{j=1}^N \int H_t(\tilde{\mathbf{s}}_{t+}) \frac{f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i})}{\bar{f}_M(\tilde{\mathbf{s}}_{t+})} f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,k}) d\tilde{\mathbf{s}}_{t+} \\ &= \frac{1}{N} \sum_{j=1}^N \int H_t(\tilde{\mathbf{s}}_{t+}) \frac{f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i})}{\bar{f}_M(\tilde{\mathbf{s}}_{t+})} \frac{1}{M} \sum_{k=1}^M f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,k}) d\tilde{\mathbf{s}}_{t+} \\ &\stackrel{(*)}{=} \frac{1}{N} \sum_{j=1}^N \int H_t(\tilde{\mathbf{s}}_{t+}) f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i}) d\tilde{\mathbf{s}}_{t+} = \frac{1}{N} \sum_{j=1}^N \mathbb{E} [H_t(\tilde{\mathbf{S}}_{t+}) | \mathbf{S}_{t,i}] \\ &= \mathbb{E} [H_t(\tilde{\mathbf{S}}_{t+}) | \mathbf{S}_{t,i}] = \Delta_t(\mathbf{S}_{t,i}), \end{aligned}$$

where (\*) holds by the definition of  $\bar{f}_M(\tilde{\mathbf{s}}_{t+})$  in (13).

Also, the likelihood ratio for the MLR estimator is always bounded above by  $M$ , because

$$\frac{f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i})}{\bar{f}_M(\tilde{\mathbf{s}}_{t+})} = \frac{f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i})}{\frac{1}{M} \sum_{j=1}^M f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,j})} \leq \frac{f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i})}{f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i})} = \frac{f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i})}{\frac{1}{M} \sum_{i=1}^M f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,i})} = M.$$

The MLR produces more stable estimates than the LR and ILR estimators because the mixture likelihood ratio is less extreme, as seen in the following example. Suppose  $f_0(x)$  and  $f_3(x)$  are the pdf's for two normal distributions, each with standard deviation equal to 1, and with means of 0 and 3 respectively. The different means can be viewed as the given outer scenarios so  $f_0$  and  $f_3$  are used to generate inner samples. Suppose a sample 3.5 is drawn from the sampling distribution  $f_0$ , and is reused in the target distribution  $f_3$ . The individual likelihood ratio is  $\frac{f_3(3.5)}{f_0(3.5)} = 403.43$ , which is extremely high; this likelihood ratio is in fact unbounded as the sample value approaches infinity. The mixture likelihood ratio, in contrast, is  $\frac{f_3(3.5)}{\frac{1}{2}[f_0(3.5)+f_3(3.5)]} = 1.995$  which is a much more reasonable value.

### 3.1.1. LR estimators with Markov state processes

When reusing simulation outputs using the ILR estimator, or the MLR estimator, the simulation outputs  $H_t(\tilde{\mathbf{S}}_{t+,kj})$  only need to be computed once, and then can be reused multiple times for different target scenarios. When the likelihood ratio calculation is faster than computing a new inner simulation output, which is often the case in practical applications, the computational saving is significant.

In general, both the individual likelihood ratio and the mixture likelihood ratio require the conditional pdf to be calculated for the whole inner sample path beyond time  $t$ ,  $\tilde{\mathbf{S}}_{t+,kj}$ , given the whole outer scenario up to time  $t$ ,  $\mathbf{S}_{t,i}$ . This calculation is simplified to a one-step transition probability  $f(\tilde{\mathbf{S}}_{t+1,kj} | \mathbf{S}_{t,i})$  when the state process is Markov, which greatly reduces the computations needed in reusing simulation outputs. The simplification as summarized in Proposition 3.1.

**Proposition 3.1.** *If Assumption 3.1 holds, and the concatenated stochastic process  $(\mathbf{S}_t, \tilde{\mathbf{S}}_{t+})$  is Markov, then the LR estimator in equation (11) can be written as*

$$\widehat{\Delta}_{t,k}^{\text{LR}}(\mathbf{S}_{t,i}) = \frac{1}{N} \sum_{j=1}^N H_t(\tilde{\mathbf{S}}_{t+,kj}) \frac{f(\tilde{\mathbf{S}}_{t+1,kj} | \mathbf{S}_{t,i})}{f(\tilde{\mathbf{S}}_{t+1,kj} | \mathbf{S}_{t,k})}, \text{ where } \tilde{\mathbf{S}}_{t+,kj} \stackrel{i.i.d.}{\sim} f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,k}) \quad \forall j = 1, \dots, N. \tag{14}$$

Moreover, the MLR estimator in equation (12) can be written as

$$\widehat{\Delta}_t^{\text{MLR}}(\mathbf{S}_{t,i}) = \frac{1}{MN} \sum_{k=1}^M \sum_{j=1}^N H_t(\tilde{\mathbf{S}}_{t+,kj}) \frac{f(\tilde{\mathbf{S}}_{t+1,kj} | \mathbf{S}_{t,i})}{\bar{f}_M(\tilde{\mathbf{S}}_{t+1,kj})},$$

where  $\bar{f}_M(\tilde{\mathbf{S}}_{t+1}) = \frac{1}{M} \sum_{i=1}^M f(\tilde{\mathbf{S}}_{t+1} | \mathbf{S}_{t,i})$  and  $\tilde{\mathbf{S}}_{t+,kj} \stackrel{i.i.d.}{\sim} f(\tilde{\mathbf{s}}_{t+} | \mathbf{S}_{t,k})$  for all  $j = 1, \dots, N$ , and for all  $k = 1, \dots, M$ .

**Proof.** By the Markov property of  $(\mathbf{S}_t, \tilde{\mathbf{S}}_{t+})$ , for any scenario  $\mathbf{S}_t$  the conditional pdf  $f(\tilde{\mathbf{S}}_{t+}|\mathbf{S}_t)$  can be written as

$$f(\tilde{\mathbf{S}}_{t+}|\mathbf{S}_t) = f(\tilde{\mathbf{S}}_{t+}|\mathbf{S}_t) = f(\tilde{\mathbf{S}}_{t+1}|\mathbf{S}_t, \tilde{\mathbf{S}}_{t+1}) \cdot f(\tilde{\mathbf{S}}_{t+1}|\mathbf{S}_t) = f(\tilde{\mathbf{S}}_{t+1}|\tilde{\mathbf{S}}_{t+1}) \cdot f(\tilde{\mathbf{S}}_{t+1}|\mathbf{S}_t),$$

where the second equality holds by the Bayes's theorem.

Then, the likelihood ratio in (14) can be simplified as

$$\frac{f(\tilde{\mathbf{S}}_{t+,kj}|\mathbf{S}_{t,i})}{f(\tilde{\mathbf{S}}_{t+,kj}|\mathbf{S}_{t,k})} = \frac{f(\tilde{\mathbf{S}}_{t+1,kj}|\tilde{\mathbf{S}}_{t+1,kj}) \cdot f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,i})}{f(\tilde{\mathbf{S}}_{t+1,kj}|\tilde{\mathbf{S}}_{t+1,kj}) \cdot f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,k})} = \frac{f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,i})}{f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,k})}.$$

For the MLR estimator, its mixture likelihood ratio can be simplified as

$$\begin{aligned} \frac{f(\tilde{\mathbf{S}}_{t+,kj}|\mathbf{S}_{t,i})}{\tilde{f}_M(\tilde{\mathbf{S}}_{t+,kj})} &= \frac{f(\tilde{\mathbf{S}}_{t+,kj}|\mathbf{S}_{t,i})}{\frac{1}{M} \sum_{i'=1}^M f(\tilde{\mathbf{S}}_{t+,kj}|\mathbf{S}_{t,i'})} \\ &= \frac{f(\tilde{\mathbf{S}}_{t+1,kj}|\tilde{\mathbf{S}}_{t+1,kj}) \cdot f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,i})}{\frac{1}{M} \sum_{i'=1}^M f(\tilde{\mathbf{S}}_{t+1,kj}|\tilde{\mathbf{S}}_{t+1,kj}) \cdot f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,i'})} \\ &= \frac{f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,i})}{\frac{1}{M} \sum_{i'=1}^M f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,i'})} = \frac{f(\tilde{\mathbf{S}}_{t+1,kj}|\mathbf{S}_{t,i})}{\tilde{f}_M(\tilde{\mathbf{S}}_{t+1,kj})}, \end{aligned} \tag{15}$$

as desired.  $\square$

Proposition 3.1 shows that even though the entire inner simulation path  $\tilde{\mathbf{S}}_{t+,kj}$  is simulated and used to calculate the simulation output  $H_t(\tilde{\mathbf{S}}_{t+,kj})$ , we only need to calculate the likelihood ratio for the one-step transition density from time  $t$  to  $t + 1$ . This means that the likelihood ratio calculation can be very efficient even for embedded options with complex path-dependent payoffs.

### 3.2. Two-stage nested simulation procedure

In Section 2 we identified two main tasks in tail risk estimation: (1) classification of tail scenarios and (2) estimation of tail losses. Following Broadie et al. (2011) and Dang et al. (2020), we use a two-stage procedure under which the first step focuses on Task (1) and the second on Task (2).

In our numerical experiments, we assume that the simulation budget is a multiple of the number of scenarios  $M$  and the number of time steps  $T$ , that is,  $\Gamma = TMN$  for some  $N$ . This setting allows us to compare our procedure with a standard multi-period nested simulation with  $N$  inner simulation outputs in each outer scenario at each time.

The two-stage procedure is as follows:

#### Stage 1: Classification of tail scenarios

In this stage, we use a fraction of the simulation budget, say  $TMN_1 < \Gamma$  for some  $N_1$ , to identify a set of tail scenarios that are highly likely to belong to the true tail scenario set,  $\mathcal{T}_\alpha$ . The number  $N_1 \leq N$  is a design parameter selected by the user.

(1.A) For each time- $t$  scenario  $\mathbf{S}_{t,i}$ ,  $i = 1, \dots, M$  and  $t = 0, \dots, T - 1$ , simulate  $N_1$  inner sample paths and compute the inner simulation outputs. There are  $MN_1$  inner simulation outputs at each time  $t = 0, \dots, T - 1$ .

(1.B) Use the MLR estimator (12) to estimate  $\Delta_t(\mathbf{S}_{t,i})$ , for all  $i = 1, \dots, M$  and  $t = 0, \dots, T - 1$ . Denote these first stage estimators by  $\hat{\Delta}_t^{\text{MLR}_1}(\mathbf{S}_{t,i})$ . The same  $MN_1$  inner samples are reused in all  $M$  target scenarios, but weighted by different likelihood ratios.

(1.C) Use the estimates  $\hat{\Delta}_t^{\text{MLR}_1}(\mathbf{S}_{t,i})$  in Step (1.B) to estimate the loss  $L_i$  by

$$\hat{L}_i^{\text{MLR}_1} = \sum_{t=0}^{T-1} \hat{\Delta}_t^{\text{MLR}_1}(\mathbf{S}_{t,i}) [e^{-rt} S_{t,i} - e^{-r(t+1)} S_{t+1,i}] + V_T(\mathbf{S}_{T,i}), \quad i = 1, \dots, M.$$

(1.D) Sort the estimated losses in Step (1.C),  $\hat{L}_{(1)}^{\text{MLR}_1} < \hat{L}_{(2)}^{\text{MLR}_1} < \dots < \hat{L}_{(M)}^{\text{MLR}_1}$  (assume no ties for simplicity). Identify the set of scenarios with the  $M^h$  highest estimated losses.

$$\hat{\mathcal{T}}_\alpha^{\text{HL}} = \left\{ \mathbf{S}_{T,i} : \hat{L}_i^{\text{MLR}_1} > \hat{L}_{(M-M^h)}^{\text{MLR}_1} \right\}.$$

The number  $M^h \geq (1 - \alpha)M$  is a design parameter selected by the user. The set  $\hat{\mathcal{T}}_\alpha^{\text{HL}}$  is called the *highly likely tail scenario set*.

#### Stage 2: Estimation of tail losses

In this stage, we concentrate the remaining simulation budget on the scenarios in  $\hat{\mathcal{T}}_\alpha^{\text{HL}}$ . Stage 1 uses  $TMN_1$  simulation outputs, so the remaining simulation budget in Stage 2 is  $\Gamma - TMN_1 = TM(N - N_1)$ . Therefore, each of the  $M^h$  highly likely tail scenarios generates approximately  $N_2 = \lceil M(N - N_1)/M^h \rceil$  simulation outputs at each time  $t = 0, \dots, T - 1$ .

(2.A) For each scenario  $\mathbf{S}_{T,i} \in \hat{\mathcal{T}}_\alpha^{\text{HL}}$ , simulate  $N_2$  inner sample paths at  $t = 0, 1, \dots, T - 1$ , and compute the inner simulation outputs. There are  $M^h N_2$  newly simulated outputs at each time  $t = 0, \dots, T - 1$ .

(2.B) Use the MLR estimator  $\hat{\Delta}_t^{\text{MLR}}(\mathbf{S}_{t,i})$  to estimate  $\Delta_t(\mathbf{S}_{t,i})$  for all  $\mathbf{S}_{T,i} \in \hat{\mathcal{T}}_\alpha^{\text{HL}}$  and for all  $t = 0, \dots, T - 1$ . The outputs simulated for  $\mathbf{S}_{t,i} \in \hat{\mathcal{T}}_\alpha^{\text{HL}}$  in Stage 1 are also included, so there are  $M^h(N_1 + N_2)$  inner simulation outputs in total that are reused by the MLR estimators.

(2.C) Use the estimates  $\hat{\Delta}_t^{\text{MLR}}(\mathbf{S}_{t,i})$ ,  $t = 0, 1, \dots, T - 1$  in Step (2.B) to estimate the losses  $\hat{L}_i^{\text{MLR}} = \hat{L}^{\text{MLR}}(\mathbf{S}_{T,i})$  for all  $\mathbf{S}_{T,i} \in \hat{\mathcal{T}}_\alpha^{\text{HL}}$ .

(2.D) Sort the estimated losses in Step (2.C),  $\widehat{L}_{(1)}^{\text{MLR}} < \widehat{L}_{(2)}^{\text{MLR}} < \dots < \widehat{L}_{(M^h)}^{\text{MLR}}$ , then estimate the  $\alpha$ -CTE by

$$\widehat{\text{CTE}}_{\alpha}^{\text{MLR}} = \frac{1}{(1-\alpha)M} \sum_{i=M^h-(1-\alpha)M}^{M^h} \widehat{L}_{(i)}^{\text{MLR}} = \frac{1}{(1-\alpha)M} \sum_{\mathbf{S}_{T,i} \in \widehat{\mathcal{T}}_{\alpha}^{\text{MLR}}} \widehat{L}_i^{\text{MLR}}, \tag{16}$$

where the set of MLR tail scenarios is

$$\widehat{\mathcal{T}}_{\alpha}^{\text{MLR}} = \{\mathbf{S}_{T,i} \in \widehat{\mathcal{T}}_{\alpha}^{\text{HL}} : \widehat{L}_i^{\text{MLR}} > \widehat{L}_{(M^h-(1-\alpha)M)}^{\text{MLR}}\}.$$

The MLR tail scenarios,  $\widehat{\mathcal{T}}_{\alpha}^{\text{MLR}}$ , are the  $(1-\alpha)M$  scenarios generating the largest estimated losses, according to the Stage 2 MLR estimators, from the  $M^h$  scenarios in  $\mathbf{S}_{T,i} \in \widehat{\mathcal{T}}_{\alpha}^{\text{HL}}$ .

We see that Steps (2.A)–(2.C) are similar to Steps (1.A)–(1.C), but the simulation and estimation involve only on the highly likely tail scenarios,  $\widehat{\mathcal{T}}_{\alpha}^{\text{HL}}$ .

Both design parameters  $N_1$  and  $M^h$  affect the classification of tail scenarios and estimation of tail losses. The optimal selection of these parameters will be considered in future studies. We find in our experiments that judicious choices of the design parameters  $N_1$  and  $M^h$  will improve performance. Here we provide some guidelines based on our experience:

- We recommend that  $N_1$ , which is the number of inner simulations per outer simulation in Stage 1, should be small, e.g., 1 or 2. In Stage 1, the goal is to identify the highly likely tail scenarios. We can afford coarse estimates for the values of the scenario losses, as long as their relative rankings are similar to the rankings of the true losses. The MLR estimator reuses the inner sample paths from all  $M$  scenarios, thus every  $\Delta_t^{\text{MLR}}(\mathbf{S}_t)$  is estimated using  $MN_1$  inner sample paths; so a small  $N_1$  suffices for our purpose in Stage 1. Moreover, the smaller  $N_1$  is, the larger the remaining simulation budget is for more accurate estimation in Stage 2.
- We recommend that  $M^h$ , which is the number of highly likely tail scenarios in Stage 1, should be the number of true tail scenarios plus a safety margin. Due to simulation noise, the rankings of the MLR loss estimates will be different from the rankings of the true losses. Ideally, we would like  $\mathcal{T}_{\alpha} \subseteq \widehat{\mathcal{T}}_{\alpha}^{\text{HL}}$ . Increasing  $M^h$  increases the probability of this, but if  $M^h$  is too large, the Stage 2 simulation budget will be smaller. The appropriate safety margin varies for different applications. We find in our experiments that a margin between 5% $M$  and 15% $M$  strikes a good balance between ensuring a good coverage of the true tail scenarios in Stage 1, and leaving sufficient simulation budget for Stage 2.

Note that  $M^h$  and  $N_1$  are connected – a smaller value for  $N_1$  will give a rougher first estimate of the losses for each scenario, which may necessitate a larger value of  $M^h$ , for example.

#### 4. Applying the two-stage MLR approach to a GMWB

##### 4.1. Financial modeling and dynamic hedging for GMWB

The policyholder of a GMWB option may periodically withdraw a guaranteed amount from their sub-account, until a prescribed time or until their death. The withdrawals often begin at a specified date (typically 5–7 years) after the initial premium investment.

The guaranteed withdrawal amount is based on a specified percentage of the guarantee base. The policyholder is entitled to withdraw this amount even if their sub-account is entirely depleted. The guarantee base is usually set at the higher watermark of the sub-account value.

Consider a GMWB with ratchet that expires in  $T$  months. We denote the sub-account value at time  $t$  by  $F_t$ . We assume that the account is invested in a stock index whose time  $t$  value is denoted by  $S_t$ . For simplicity, we assume that the withdrawal benefit starts immediately after the contract commences.

Let  $G_t$  be the ratcheted guarantee base at time  $t$ . The contract offers a guaranteed periodic withdrawal benefit in the amount of  $I_t = \gamma G_t$  for some fixed  $\gamma$ . We assume the GMWB is issued to a policyholder age  $x$  at  $t = 0$ . Let  ${}_s p_x$  denote the probability of this policyholder surviving all decrements, including mortality and lapse, up to time  $s$ , in months, and we let  $p_{x,t}$  denote the probability of survival from time  $t$  to time  $t + 1$ , in months. For monthly withdrawals in a 20-year contract, we need to consider the evolution of  $F_t$ ,  $S_t$ , and  $G_t$ , for  $t = 0, 1, 2, \dots, T$ . At  $t = 0$ , we assume that the whole fund is invested in the stock index and the guarantee base is set to the fund value, so

$$F_0 = S_0 = G_0.$$

Also, set  $I_0 = 0$  as the first withdrawal starts at time 1.

At each withdrawal date, the follow events take place in order:

- (1) The fund value from  $t - 1$ , after the withdrawal at that time, changes according to the growth of the underlying stock, deduction of management fee, and reduction proportional to mortality and lapse decrements.
- (2) The guarantee base ratchets up if the fund value exceeds the previous guarantee base, reduced proportionally for decrements.
- (3) The fund value is reduced by the amount of the withdrawal benefit, subject to a minimum value of 0.

Mathematically, at  $t = 1, \dots, T$ , we have

$$F_t = \max \left( (F_{t-1} - I_{t-1}) \frac{S_t}{S_{t-1}} e^{-\eta_g} \cdot p_{x,t-1}, 0 \right) = (F_{t-1} - I_{t-1})^+ e^{Rt} e^{-\eta_g} \cdot p_{x,t-1}, \tag{17}$$

$$\text{where } R_t = \ln \frac{S_t}{S_{t-1}}$$

$$G_t = \max(G_{t-1} \cdot p_{x+t-1}, F_t), \tag{18}$$

$$I_t = \gamma G_t. \tag{19}$$

Here  $R_t$  is the log-return of the stock index at  $t$ , and  $\eta_g$  is the gross management fee that is deducted from the sub-account each month. Part of the gross management fee covers the expenses incurred by the insurer and the other part is the income for the insurer. Denote the net-of-expenses rate of management fee by  $\eta_n$ ; thus the insurer's income at time  $t$  is  $F_t(e^{\eta_n} - 1)$ .

We see from (17)–(19) that the status of a GMWB at any time  $t$  can be summarized by the triplet  $(S_t, F_t, G_t)$ , so this triplet is the state variable. We note that the evolution of the triplet is driven only by the stochasticity of the stock price process  $S_t$ . We denote the state variables up to time  $t$  by  $(\mathbf{S}_t, \mathbf{F}_t, \mathbf{G}_t)$  and the inner simulated state variables from time  $t$  to  $T$  by  $(\tilde{\mathbf{S}}_{t+}, \tilde{\mathbf{F}}_{t+}, \tilde{\mathbf{G}}_{t+})$ .

The insurer's liability in a GMWB contract is the present value of all the withdrawal benefits paid after the depletion of the sub-account, offset by the present value of all the net management fees,  $F_t(e^{\eta_n} - 1)$ , collected as income. At any time  $t$ , the insurer is interested in hedging the future liability (beyond time  $t$ ), which has value at  $t$  of

$$v_{t+} = \sum_{s=t+1}^T e^{-r(s-t)} [(I_s - F_s)^+ - F_s(e^{\eta_n} - 1)]. \tag{20}$$

Consider a discrete time delta hedging program for this GMWB contract. The insurer sets up an initial hedge portfolio at time 0 then rebalances it at times  $t = 1, \dots, T$  to offset the future liabilities, based on a delta hedge. The time  $t$  hedge portfolio consists of  $\Delta_t$  units of the underlying stock  $S_t$ , where  $\Delta_t$  is the derivative of the GMWB liability value at  $t$ , with respect to the time  $t$  stock price  $S_t$ . This derivative can be estimated as follows.

1. If  $F_t \leq \gamma G_t$ , then  $\Delta_t = 0$  because the sub-account would be depleted after the withdrawal at time  $t$  so the liability value does not depend on  $S_t$ . In this case, we set  $\Delta_t = 0$  without any inner simulation. In implementation, these scenarios are flagged as they do not reuse any output and do not create any simulation output to be reused by other scenarios.
2. If  $F_t > \gamma G_t$ , then  $\Delta_t = \Delta_t(\mathbf{S}_t, \mathbf{F}_t, \mathbf{G}_t)$  is estimated using inner simulation. Following Glasserman (2013) and Cathcart et al. (2015), based on equation (20) the pathwise estimator of  $\Delta_t(\mathbf{S}_t, \mathbf{F}_t, \mathbf{G}_t)$ , based on a single inner simulation path, is given by

$$H_t(\tilde{\mathbf{S}}_{t+}, \tilde{\mathbf{F}}_{t+}, \tilde{\mathbf{G}}_{t+}) = \sum_{s=t+1}^T e^{-r(s-t)} \left[ \mathbb{1}\{\tilde{I}_s > \tilde{F}_s\} \cdot \left( \frac{d\tilde{I}_s}{dS_t} - \frac{d\tilde{F}_s}{dS_t} \right) - \frac{d\tilde{F}_s}{dS_t} (e^{\eta_n} - 1) \right]. \tag{21}$$

Using equations (17)–(19), the sensitivities

$$\frac{d\tilde{F}_s}{dS_t}, \quad \frac{d\tilde{G}_s}{dS_t}, \quad \text{and} \quad \frac{d\tilde{I}_s}{dS_t}$$

are calculated recursively for  $s = t+1, t+2, \dots, T$  as

$$\frac{d\tilde{F}_s}{dS_t} = \mathbb{1}\{\tilde{I}_{s-1} < \tilde{F}_{s-1}\} \cdot \left( \frac{d\tilde{F}_{s-1}}{dS_t} - \frac{d\tilde{I}_{s-1}}{dS_t} \right) e^{\tilde{R}_s} e^{-\eta_g} \cdot p_{x,s-1}, \text{ where } \tilde{R}_s = \ln \frac{\tilde{S}_s}{\tilde{S}_{s-1}} \tag{22}$$

$$\frac{d\tilde{G}_s}{dS_t} = \mathbb{1}\{\tilde{G}_{s-1} \cdot p_{x,s-1} < \tilde{F}_s\} \frac{d\tilde{F}_s}{dS_t} + \mathbb{1}\{G_{s-1} \cdot p_{x,s-1} \geq F_s\} \frac{d\tilde{G}_{s-1}}{dS_t}, \tag{23}$$

$$\frac{d\tilde{I}_s}{dS_t} = \gamma \frac{d\tilde{G}_s}{dS_t}. \tag{24}$$

The boundary conditions of these recursions are given by

$$\frac{d\tilde{F}_t}{dS_t} = \frac{F_t}{S_t}, \quad \frac{d\tilde{G}_t}{dS_t} = 0, \quad \text{and} \quad \frac{d\tilde{I}_t}{dS_t} = 0, \tag{25}$$

because

- At time  $t$ , the fund  $F_t$  has  $\frac{F_t}{S_t}$  unit of stocks so a unit change in stock price results in  $\frac{F_t}{S_t}$  units of change in fund value, i.e.,

$$\frac{dF_t}{dS_t} = \frac{F_t}{S_t}.$$

- Also, the inner simulation sample path at time  $t$  is initialized by setting  $(\tilde{S}_t, \tilde{F}_t, \tilde{G}_t) = (S_t, F_t, G_t)$ . Therefore

$$\frac{d\tilde{F}_t}{dS_t} = \frac{dF_t}{dS_t} = \frac{F_t}{S_t}.$$

- Given the guarantee base  $G_t$ , the guarantee base  $\tilde{G}_t = G_t$  and the withdrawal amount  $\tilde{I}_t = \gamma G_t$  is fixed, which means that

$$\frac{d\tilde{G}_t}{dS_t} = 0 \text{ and } \frac{d\tilde{I}_t}{dS_t} = 0.$$

Note that inner simulation is only required when  $F_t > \gamma G_t$ , in which case the indicator function in equation (22) is equal to 1. Also, based on the boundary conditions (25), provided that  $F_t > \gamma G_t$ , we have

$$\frac{d\tilde{F}_{t+1}}{dS_t} = \left( \frac{d\tilde{F}_t}{dS_t} - \frac{d\tilde{I}_t}{dS_t} \right) e^{\tilde{R}_{t+1}} e^{-\eta_g} \cdot p_{x,t} = \frac{F_t}{S_t} e^{\tilde{R}_{t+1}} e^{-\eta_g} \cdot p_{x,t}, \tag{26}$$

where  $\tilde{R}_{t+1} = \ln \frac{\tilde{S}_{t+1}}{\tilde{S}_t} = \ln \frac{\tilde{S}_{t+1}}{S_t}$ .

In summary, in a standard nested simulation, for any scenario  $(S_{t,i}, F_{t,i}, G_{t,i})$ ,  $i = 1, \dots, M$ , the time  $t$  pathwise delta estimator for GMWB is given by

$$\hat{\Delta}_t^{NS}(S_{t,i}, F_{t,i}, G_{t,i}) = \begin{cases} 0, & \text{if } F_{t,i} \leq \gamma G_{t,i}, \\ \frac{1}{N} \sum_{j=1}^N H_t(\tilde{S}_{t+,ij}, \tilde{F}_{t+,ij}, \tilde{G}_{t+,ij}), & \text{if } F_{t,i} > \gamma G_{t,i}, \end{cases} \tag{27}$$

where  $\tilde{S}_{t+,ij} \stackrel{i.i.d.}{\sim} f(\tilde{s}_{t+}|S_{t,i})$  for all  $j = 1, \dots, N$ . Then, in a standard multi-period nested simulation procedure, for the  $i$ th scenario  $S_{T,i}$ , the discrete time hedging loss is estimated as

$$\hat{L}_i^{NS} = \sum_{t=0}^{T-1} \hat{\Delta}_t^{NS}(S_{t,i}, F_{t,i}, G_{t,i}) [e^{-rt} S_{t,i} - e^{-r(t+1)} S_{t+1,i}] + v_0(S_{T,i}, F_{T,i}, G_{T,i}), \tag{28}$$

where  $v_0(S_{T,i}, F_{T,i}, G_{T,i}) = \sum_{t=1}^T e^{-rt} [(I_t - F_t)^+ - F_t(e^{\eta_n} - 1)]$  is the value at the policy inception of the GMWB liability.

Note the similarity between equation (28) and equation (9), where the former specifically considers GMWB’s liability and estimates the deltas by the Monte Carlo estimator (27).

#### 4.2. Adapting the likelihood ratio method to the GMWB loss

To conform with the evolution of the GMWB state variables, we need to adapt the likelihood ratio method before applying it to the delta hedging of GMWB’s liability.

Recall from Section 3.1 that when reusing inner simulation outputs from sampling scenario  $k$  to a target scenario  $i$ , the inner sample paths from scenario  $k$  are concatenated with the outer path (up to time  $t$ ) from scenario  $i$ , to form a notional path:

$$\underbrace{(S_{0,i}, F_{0,i}, G_{0,i}), \dots, (S_{t,i}, F_{t,i}, G_{t,i})}_{(S_{t,i}, F_{t,i}, G_{t,i})}, \overbrace{(\tilde{S}_{t+1,kj}, \tilde{F}_{t+1,kj}, \tilde{G}_{t+1,kj}), \dots, (\tilde{S}_{T,kj}, \tilde{F}_{T,kj}, \tilde{G}_{T,kj})}^{\text{concatenated time-}(t+1)\text{ inner simulation step}} \underbrace{(\tilde{S}_{t+,kj}, \tilde{F}_{t+,kj}, \tilde{G}_{t+,kj})}_{(\tilde{S}_{t+,kj}, \tilde{F}_{t+,kj}, \tilde{G}_{t+,kj})}$$

If this notional path is valid, by which we mean that it represents a feasible evolution of the state variables, the  $(kj)$ -th simulation output is reused as the output of this notional path, after reweighting by the appropriate likelihood. However, some of these notional paths are invalid, as they do not satisfy the relationships among  $S_t$ ,  $F_t$ , and  $G_t$  implied by equations (17)–(19). We make adaptations to these paths and to the corresponding simulation outputs so they can still be reused.

From Equation (18), we see that  $G_{t+1} = \max\{F_s \cdot {}_{t-s+1}p_{x,s} : s = 0, \dots, t + 1\}$  so the ratcheted guarantee base without decrement can never decrease. Also, if the guarantee base, without decrements, is increased at time  $t + 1$ , then it must be the case that the fund value  $F_{t+1}$  has reached a high water mark. Mathematically,  $G_{t+1} > G_t \cdot p_{x,t}$  only if  $G_{t+1} = F_{t+1}$ . A concatenated notional path can violate these relationships in two ways:

- (1)  $\tilde{G}_{t+1,kj} < G_{t,i} \cdot p_{x,t}$ : The generation of the  $(kj)$ -th inner sample path is conditioned on the sampling scenario  $k$ , so its guarantee base, without decrement, is non-decreasing within that scenario, i.e.,  $\tilde{G}_{t+1,kj} \geq G_{t,k} \cdot p_{x,t}$ . However, when concatenated with the target scenario  $i$ , one may have  $\tilde{G}_{t+1,kj} < G_{t,i} \cdot p_{x,t}$ , which forms an invalid notional path.
- (2)  $\tilde{G}_{t+1,kj} > G_{t,i} \cdot p_{x,t}$  but  $\tilde{G}_{t+1,kj} \neq \tilde{F}_{t+1,kj}$ : Though  $\tilde{G}_{t+1,kj}$  is the running maximum of fund values up to time  $t + 1$  for the sampling scenario  $k$ , it may not be the running maximum of the target scenario  $S_{T,i}$  fund value up to time  $t + 1$ .

We do not want to remove all the invalid notional paths, as this is a wasteful use of simulation budget. Instead, we adjust the inner sample paths so that, after the adjustment, the concatenated notional paths are valid. We also adjust the corresponding inner simulation outputs accordingly so they can be reused in different target scenarios.

Consider a fixed time  $t = 1, \dots, T - 1$ , a target scenario  $S_{T,i}$ , such that  $F_{t,i} > I_{t,i}$ , and a sampling scenario  $S_{T,k}$ , that satisfies  $F_{t,k} > I_{t,k}$ . As discussed in Section 4.1, at time  $t$ , scenario  $S_{T,i}$  requires inner simulation and scenario  $S_{T,k}$  has inner simulation outputs to be reused. In what follows, we adjust the  $(j)$ -th inner sample path from  $S_{T,k}$ , and its simulation output, for reuse in the target scenario  $S_{T,i}$

Consider the following adjusted  $j$ -th inner sample path from the sampling scenario,  $\mathbf{S}_{T,k}$ .

$$\begin{pmatrix} \tilde{\mathbf{S}}_{t+,kj}^{\text{adj}} \\ \tilde{\mathbf{F}}_{t+,kj}^{\text{adj}} \\ \tilde{\mathbf{G}}_{t+,kj}^{\text{adj}} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{S}}_{t+,kj} \cdot \frac{S_{t,i}}{S_{t,k}} \cdot \frac{G_{t,i}}{G_{t,k}} \cdot \frac{F_{t,k} - I_{t,k}}{F_{t,i} - I_{t,i}} \\ \tilde{\mathbf{F}}_{t+,kj} \cdot \frac{G_{t,i}}{G_{t,k}} \\ \tilde{\mathbf{G}}_{t+,kj} \cdot \frac{G_{t,i}}{G_{t,k}} \end{pmatrix}. \tag{29}$$

We claim that the adjusted notional path  $\left( (\mathbf{S}_{t,i}, \mathbf{F}_{t,i}, \mathbf{G}_{t,i}), (\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}}, \tilde{\mathbf{F}}_{t+,kj}^{\text{adj}}, \tilde{\mathbf{G}}_{t+,kj}^{\text{adj}}) \right)$  is a valid evolution of GMWB state variables. First, we re-examine the two aforementioned violations:

(I) Since  $\tilde{\mathbf{G}}_{t+,kj}$  is the  $j$ th inner sample path from the  $k$ th scenario, the guarantee base without decrement is non-decreasing so  $\tilde{\mathbf{G}}_{t+1,kj} \geq G_{t,k} \cdot p_{x,t}$ . By construction, in equation (29),

$$\tilde{\mathbf{G}}_{t+1,kj}^{\text{adj}} = \tilde{\mathbf{G}}_{t+1,kj} \cdot \frac{G_{t,i} \cdot p_{x,t}}{G_{t,k} \cdot p_{x,t}} \geq G_{t,i} \cdot p_{x,t}$$

so the first violation no longer occurs.

(II) Consider any adjusted notional path where  $\tilde{\mathbf{G}}_{t+1,kj}^{\text{adj}} > G_{t,i} \cdot p_{x,t}$ . By equation (29), this means that  $\tilde{\mathbf{G}}_{t+1,kj} \cdot \frac{G_{t,i} \cdot p_{x,t}}{G_{t,k} \cdot p_{x,t}} > G_{t,i} \cdot p_{x,t}$  and so  $\tilde{\mathbf{G}}_{t+1,kj} > G_{t,k} \cdot p_{x,t}$ . Since  $\tilde{\mathbf{G}}_{t+,kj}$  is from an inner sample path of the  $k$ th scenario, if  $\tilde{\mathbf{G}}_{t+1,kj} > G_{t,k} \cdot p_{x,t}$  then it must be that  $\tilde{\mathbf{G}}_{t+1,kj} = \tilde{\mathbf{F}}_{t+1,kj}$ . Then, by (29), in this case

$$\tilde{\mathbf{G}}_{t+1,kj}^{\text{adj}} = \tilde{\mathbf{G}}_{t+1,kj} \cdot \frac{G_{t,i}}{G_{t,k}} = \tilde{\mathbf{F}}_{t+1,kj} \cdot \frac{G_{t,i}}{G_{t,k}} = \tilde{\mathbf{F}}_{t+1,kj}^{\text{adj}}$$

so the second violation no longer occurs either.

Secondly, we show that the adjusted stock price path  $\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}}$  is aligned with the adjustments made to  $\tilde{\mathbf{F}}_{t+,kj}^{\text{adj}}$  and  $\tilde{\mathbf{G}}_{t+,kj}^{\text{adj}}$ . By equations (17) and (29), we have

$$(F_{t,i} - I_{t,i})^+ \cdot \frac{\tilde{\mathbf{S}}_{t+1,kj}^{\text{adj}}}{S_{t,i}} e^{-\eta^g} \cdot p_{x,t} \stackrel{(17)}{=} \tilde{\mathbf{F}}_{t+1,kj}^{\text{adj}} \stackrel{(29)}{=} \tilde{\mathbf{F}}_{t+1,kj} \cdot \frac{G_{t,i}}{G_{t,k}} \stackrel{(17)}{=} (F_{t,k} - I_{t,k})^+ \cdot \frac{\tilde{\mathbf{S}}_{t+1,kj}}{S_{t,k}} e^{-\eta^g} \cdot \frac{G_{t,i}}{G_{t,k}} \cdot p_{x,t}.$$

As we consider target and sampling scenarios with  $F_{t,i} > I_{t,i}$  and  $F_{t,k} > I_{t,k}$ , the adjusted stock price is

$$\tilde{\mathbf{S}}_{t+1,kj}^{\text{adj}} = \tilde{\mathbf{S}}_{t+1,kj} \cdot \frac{S_{t,i}}{S_{t,k}} \cdot \frac{G_{t,i}}{G_{t,k}} \cdot \frac{F_{t,k} - I_{t,k}}{F_{t,i} - I_{t,i}} = S_{t,i} \cdot e^{\tilde{\mathbf{R}}_{t+1,kj}^{\text{adj}}},$$

where

$$\tilde{\mathbf{R}}_{t+1,kj}^{\text{adj}} = \ln \frac{\tilde{\mathbf{S}}_{t+1,kj}^{\text{adj}}}{S_{t,i}} = \tilde{\mathbf{R}}_{t+1,kj} + \ln \left[ \frac{G_{t,i}}{G_{t,k}} \cdot \frac{F_{t,k} - I_{t,k}}{F_{t,i} - I_{t,i}} \right] \text{ and } \tilde{\mathbf{R}}_{t+1,kj} = \ln \frac{\tilde{\mathbf{S}}_{t+1,kj}}{S_{t,k}}.$$

This shows that, to align with the adjusted fund value  $\tilde{\mathbf{F}}_{t+1}^{\text{adj}}$  and guarantee base  $\tilde{\mathbf{G}}_{t+1}^{\text{adj}}$  at time  $t + 1$ , the time  $(t + 1)$  log-return is adjusted from  $\tilde{\mathbf{R}}_{t+1,kj}$  to  $\tilde{\mathbf{R}}_{t+1,kj}^{\text{adj}}$ . For subsequent time steps no adjustment is made to the log-returns, so  $\frac{\tilde{\mathbf{S}}_s^{\text{adj}}}{\tilde{\mathbf{S}}_{s-1}^{\text{adj}}} = \frac{\tilde{\mathbf{S}}_s}{\tilde{\mathbf{S}}_{s-1}}$ . Then, according to equations (17)–(19), the stock price  $\tilde{\mathbf{S}}_s^{\text{adj}}$ , the fund value  $\tilde{\mathbf{F}}_s^{\text{adj}}$ , and the guarantee based  $\tilde{\mathbf{G}}_s^{\text{adj}}$  all are the same multiple of their unadjusted value as they are at time  $t + 1$ . This justifies the proportionate change for the entire inner path in equation (29).

The adjustment (29) enables us to reuse the adjusted inner sample paths. The simulation output for an adjusted inner sample path is  $H_t(\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}}, \tilde{\mathbf{F}}_{t+,kj}^{\text{adj}}, \tilde{\mathbf{G}}_{t+,kj}^{\text{adj}})$ , where  $H_t$  is as defined in equation (21), and requires the recursions from equations (22)–(25). We show that the recursion can be circumvented by reusing the unadjusted simulation outputs  $H_t(\tilde{\mathbf{S}}_{t+,kj}, \tilde{\mathbf{F}}_{t+,kj}, \tilde{\mathbf{G}}_{t+,kj})$ . Specifically,

$$H_t(\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}}, \tilde{\mathbf{F}}_{t+,kj}^{\text{adj}}, \tilde{\mathbf{G}}_{t+,kj}^{\text{adj}}) = H_t(\tilde{\mathbf{S}}_{t+,kj}, \tilde{\mathbf{F}}_{t+,kj}, \tilde{\mathbf{G}}_{t+,kj}) \cdot \frac{e^{\tilde{\mathbf{R}}_{t+1,kj}^{\text{adj}}}}{e^{\tilde{\mathbf{R}}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}} \cdot \frac{S_{t,k}}{F_{t,k}}, \tag{30}$$

where

$$\tilde{\mathbf{R}}_{t+1,kj} = \ln \frac{\tilde{\mathbf{S}}_{t+1,kj}}{S_{t,k}} \text{ and } \tilde{\mathbf{R}}_{t+1,kj}^{\text{adj}} = \tilde{\mathbf{R}}_{t+1,kj} + \ln \left[ \frac{G_{t,i}}{G_{t,k}} \cdot \frac{F_{t,k} - I_{t,k}}{F_{t,i} - I_{t,i}} \right].$$

Firstly, the adjustment in equation (29) implies that, for any  $s = t + 1, \dots, T$ ,

$$\frac{d\tilde{\mathbf{F}}_{s,kj}^{\text{adj}}}{d\tilde{\mathbf{F}}_{s,kj}} = \frac{d\tilde{\mathbf{F}}_{t+1,kj}^{\text{adj}}}{d\tilde{\mathbf{F}}_{t+1,kj}} = \frac{G_{t,i}}{G_{t,k}}.$$

Using the chain rule, we have

$$\frac{d\tilde{F}_{s,kj}^{adj}}{dS_{t,i}} = \frac{d\tilde{F}_{s,kj}}{dS_{t,k}} \cdot \frac{d\tilde{F}_{t+1,kj}^{adj}}{dS_{t,i}} \quad (26) \quad \frac{d\tilde{F}_{s,kj}}{dS_{t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{adj}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}}, \quad \forall s = t + 1, \dots, T. \quad (31)$$

Equation (31) can be interpreted as two adjustments to the  $(kj)$ -th simulation output: one for the log-return and one for the boundary condition. Both adjustments are results of the inner path adjustment from equation (29).

Secondly, the adjustment implies that, for any  $s = t + 1, \dots, T$ ,

$$\begin{aligned} \frac{d\tilde{G}_{s,kj}^{adj}}{d\tilde{G}_{s,kj}} &= \frac{d\tilde{F}_{s,kj}^{adj}}{d\tilde{F}_{s,kj}} = \frac{G_{t,i}}{G_{t,k}} \Rightarrow \frac{d\tilde{G}_{s,kj}^{adj}}{d\tilde{F}_{s,kj}^{adj}} = \frac{d\tilde{G}_{s,kj}}{d\tilde{F}_{s,kj}} \\ \Rightarrow \frac{d\tilde{G}_{s,kj}^{adj}/dS_{t,i}}{d\tilde{F}_{s,kj}^{adj}/dS_{t,i}} &= \frac{d\tilde{G}_{s,kj}^{adj}}{d\tilde{F}_{s,kj}^{adj}} = \frac{d\tilde{G}_{s,kj}}{d\tilde{F}_{s,kj}} = \frac{d\tilde{G}_{s,kj}/dS_{t,k}}{d\tilde{F}_{s,kj}/dS_{t,k}} \\ \Rightarrow \frac{d\tilde{G}_{s,kj}^{adj}}{dS_{t,i}} &= \frac{d\tilde{G}_{s,kj}}{dS_{t,k}} \frac{d\tilde{F}_{s,kj}^{adj}}{dS_{t,i}} \quad (31) \quad \frac{d\tilde{G}_{s,kj}}{dS_{t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{adj}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}}, \quad \forall s = t + 1, \dots, T. \end{aligned} \quad (32)$$

Thirdly, the withdrawal  $I_t$  is a fixed proportion ( $\gamma$ ) of the guarantee base, regardless of the adjustment. This means that

$$\begin{aligned} \frac{d\tilde{I}_{s,kj}^{adj}}{d\tilde{G}_{s,kj}^{adj}} &= \frac{d\tilde{I}_{s,kj}}{d\tilde{G}_{s,kj}} = \gamma \\ \Rightarrow \frac{d\tilde{I}_{s,kj}^{adj}/dS_{t,i}}{d\tilde{G}_{s,kj}^{adj}/dS_{t,i}} &= \frac{d\tilde{I}_{s,kj}/dS_{t,i}}{d\tilde{G}_{s,kj}/dS_{t,i}} \\ \Rightarrow \frac{d\tilde{I}_{s,kj}^{adj}}{dS_{t,i}} &= \frac{d\tilde{I}_{s,kj}}{dS_{t,i}} \frac{d\tilde{G}_{s,kj}^{adj}}{dS_{t,i}} \quad (32) \quad \frac{d\tilde{I}_{s,kj}}{dS_{t,i}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{adj}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}}, \quad \forall s = t + 1, \dots, T. \end{aligned} \quad (33)$$

Finally, we see that equation (30) holds by plugging equations (31)–(33) into equation (21).

Let  $\Lambda_t := \{(S_{T,i}, F_{T,i}, G_{T,i}) : i = 1, \dots, M, F_{T,i} > I_{T,i}\}$  be the outer scenarios that require inner simulations at time  $t$  and let  $M_t = |\Lambda_t|$ . In light of (12), (27), and (30), for any target scenario  $(S_{T,i}, F_{T,i}, G_{T,i}) \in \Lambda_t$ , the MLR estimator for the time  $t$  delta is

$$\hat{\Delta}_t^{MLR}(S_{t,i}, F_{t,i}, G_{t,i}) = \begin{cases} 0 & \text{if } F_{t,i} \leq \gamma G_{t,i}, \\ \frac{1}{M_t N} \sum_{k:(S_{T,k}, F_{T,k}, G_{T,k}) \in \Lambda_t} \frac{e^{\tilde{R}_{t+1,kj}^{adj}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}} \cdot \frac{F_{t,k}}{S_{t,k}} & \text{if } F_{t,i} > \gamma G_{t,i}, \\ \times \sum_{j=1}^N H_t(\tilde{S}_{t+,kj}, \tilde{F}_{t+,kj}, \tilde{G}_{t+,kj}) \frac{f(\tilde{S}_{t+,kj}^{adj} | S_{t,i})}{\tilde{f}_M(\tilde{S}_{t+,kj}^{adj})} & \end{cases} \quad (34)$$

where

$$\tilde{R}_{t+1,kj} = \ln \frac{\tilde{S}_{t+1,kj}}{S_{t,k}} \quad \text{and} \quad \tilde{R}_{t+1,kj}^{adj} = \tilde{R}_{t+1,kj} + \ln \left[ \frac{G_{t,i}}{G_{t,k}} \cdot \frac{F_{t,k} - I_{t,k}}{F_{t,i} - I_{t,i}} \right].$$

The likelihood ratio  $\frac{f(\tilde{S}_{t+,kj}^{adj} | S_{t,i})}{\tilde{f}_M(\tilde{S}_{t+,kj}^{adj})}$  can be calculated using equation (15) if the state process is Markov. We provide an example of its calculation in Section 5.

We see from equation (34) that the unadjusted simulations outputs  $H_t(\tilde{S}_{t+,kj}, \tilde{F}_{t+,kj}, \tilde{G}_{t+,kj})$  are computed once and reused  $M_t$  times in different target scenarios. The adjustment adds negligible computation.

As mentioned in Section 4.1, the stochasticity of GMWB state variable  $(S_t, F_t, G_t)$  is driven entirely by the underlying stock  $S_t$ , so the likelihood ratio calculation in (34) is based on the conditional densities of inner sample paths of  $\tilde{S}_{t+}$ , given the outer scenario. The adjustment in equation (29) does not affect the Markov property of the state process. To calculate the MLR estimator for other types of VA contracts, particularly where the stochastic state variables are driven by the underlying stock price and are path-dependent, similar adaptation can be made. We show an example using a GMMB contract in Section 6.

In the following section we illustrate the two-stage procedure numerically for a fixed term GMWB. In Appendix B, we also present results of a numerical experiment using a simplified GMMB contract with a log-normal asset model.

**Table 1**  
Parameters for the regime-switching lognormal asset model in the numerical experiments.

(Monthly rate)	Real World
Risk-free Rate: $r$	0.002
Mean - Regime 1 ( $\rho = 1$ ): $\mu_1$	0.0085
Mean - Regime 2 ( $\rho = 2$ ): $\mu_2$	-0.0200
Standard Deviation - Regime 1: $\sigma_1$	0.035
Standard Deviation - Regime 2: $\sigma_2$	0.080
Transition Probability - from Regime 1: $p_{12}$	0.04
Transition Probability - from Regime 2: $p_{21}$	0.20

### 5. Numerical experiments

In this section we compare a standard, multi-period nested simulation to the two-stage procedure described above, in the context of estimating the 95%-CTE of the hedging losses for a delta-hedged GMWB contract. Our experiments show that the proposed two-stage procedure produces more accurate estimates than the standard procedure with much less runtime.

We consider a fixed term GMWB contract with the following characteristics:

- The initial fund value is  $F_0 = 1000$ , which is also the initial guarantee base  $G_0$ . The guarantee base is ratcheted monthly to the higher of the previous month's guarantee base and the current month's fund value prior to the withdrawal.
- It has a 20-year time-to-maturity and is hedged monthly, so  $T = 240$ .
- The policyholder is allowed to make monthly withdrawals of  $\gamma = 0.375\%$  of the guarantee base  $G_t$ .
- A management fee of  $\eta_g = 0.2\%$  is deducted monthly from the fund value and half of it is treated as income for the contract guarantees, that is,  $\eta_n = 0.1\%$ .
- For the purpose of comparison between the methods, we simplify the calculations by ignoring mortality, and assuming a constant risk-free interest rate,  $r$ . We also assume that withdrawals exactly meet the guaranteed minimum each month.

We assume that the underlying stock follows a regime-switching lognormal (RSLN) asset model with two regimes, with parameters as specified in Table 1 for the real world measure. The risk neutral measure is identical, except that the mean log returns are  $r - \frac{\sigma_i^2}{2}$  for regimes  $i = 1, 2$ ; the other parameters are unchanged. This is a common approach in the literature (See Bollen, 1998; Hardy, 2001; Dang et al., 2020, for example). The RSLN model is popular in practice for its ability to model volatility clustering and other asset price characteristics. It also satisfies the Markov property, allowing the use of Proposition 3.1 to simplify the likelihood ratio calculations.

The likelihood ratio in the MLR estimator in equation (34) is calculated as

$$\begin{aligned} \frac{f(\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}}|\mathbf{S}_{t,i})}{\tilde{f}_M(\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}})} &= \frac{f(\tilde{\mathbf{S}}_{t+,1,kj}|\mathbf{S}_{t,i})}{\frac{1}{M} \sum_{i'=1}^M f(\tilde{\mathbf{S}}_{t+,1,kj}|\mathbf{S}_{t,i'})} \\ &= \frac{\phi\left(\ln\left(\frac{\tilde{S}_{t+,1,kj}}{S_{t,i}}\right); r - \frac{\sigma_{\rho_{t,i'}}^2}{2}, \sigma_{\rho_{t,i}}^2\right) \cdot P[\rho_{t+,1,kj}|\rho_{t,i}]}{\frac{1}{M} \sum_{i'=1}^M \phi\left(\ln\left(\frac{\tilde{S}_{t+,1,kj}}{S_{t,i'}}\right); r - \frac{\sigma_{\rho_{t,i'}}^2}{2}, \sigma_{\rho_{t,i'}}^2\right) \cdot P[\rho_{t+,1,kj}|\rho_{t,i'}]} \end{aligned}$$

where  $\phi(x; \mu; \sigma^2)$  is the probability density function of a normal distribution with mean  $\mu$  and variance  $\sigma^2$ .

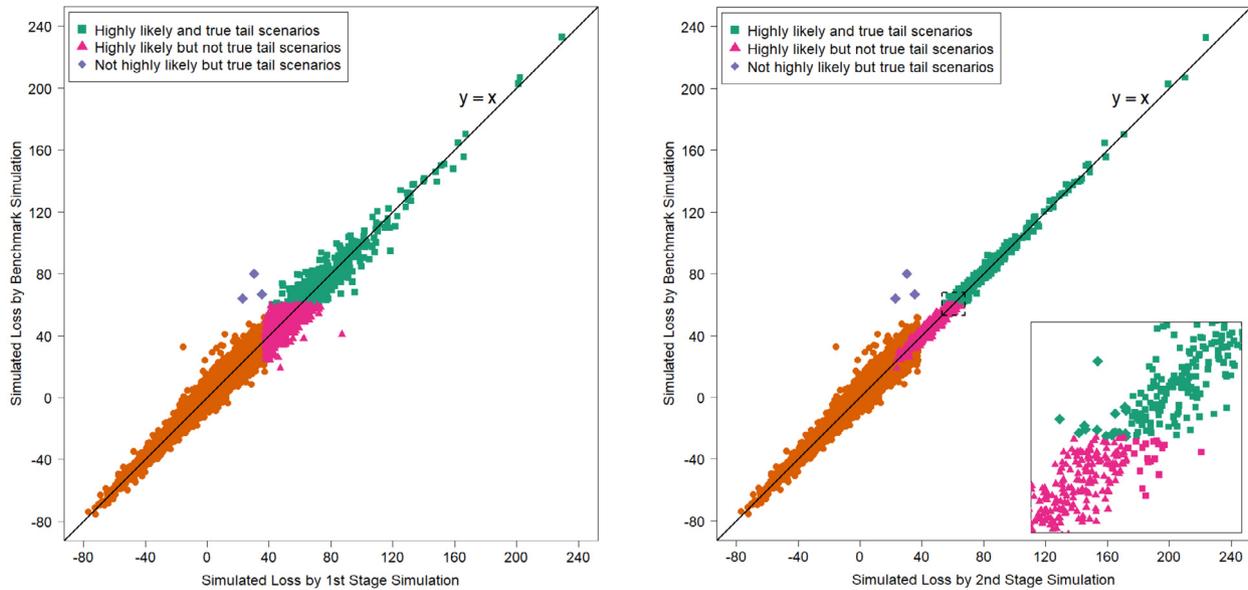
We first simulate  $M = 10,000$  stock paths  $\mathbf{S}_{T,i}$ ,  $i = 1, \dots, M$ , and use them to compute the fund value paths  $\mathbf{F}_{T,i}$  and guarantee base paths  $\mathbf{G}_{T,i}$ . All experiments in this section will use the same set of scenarios. The objective is to estimate the tail risk measure as accurately and efficiently as possible for this fixed set of scenarios. This requires identifying the 500 scenarios generating the largest losses, and accurately evaluating the losses for these scenarios.

As the hedging loss for the GMWB contract under the RSLN model cannot be calculated analytically, we run a large scale standard nested simulation with  $N = 10,000$  inner simulations at each time step of each scenario, to obtain accurate estimates for the hedging loss for each scenario. These accurate estimates are used as benchmarks to assess the accuracy of different simulation procedures. We are interested in estimating the 95%-CTE of these benchmark losses, which is the average of the largest 500 benchmark losses corresponding to the true tail scenarios.

We first provide a holistic view of the performance of the two-stage procedure. Though not all scenarios require inner simulation, we assign the same number of inner sample paths (some unused) to all scenarios. Consider a two-stage procedure with  $N_1 = 2$  inner sample paths per sampling scenario (for a total of 20,000 inner simulation outputs reused in each target scenarios) in Stage 1, and  $M^h = 1,500$  highly likely tail scenarios (that is, a 10% $M$  safety margin).

In Stage 1, the MLR estimator reuses 20,000 simulation outputs to estimate the loss in each scenario. Fig. 3a depicts the estimated losses in Stage 1 versus the benchmark losses. We see that the MLR estimates of the losses are close to the benchmark losses, as the points in Fig. 3a are near the 45-degree line. Moreover, 497 of the 500 true tail scenarios are included in the  $M^h = 1,500$  highly likely tail scenarios.

In Stage 2, each of the 1,500 highly likely tail scenarios are assigned an additional  $N_2 = 80$  inner simulations, so the MLR estimator reuses  $(N_1 + N_2) \cdot M^h = 123,000$  simulation outputs to estimate the loss in each scenario. Fig. 3b shows the estimated losses in Stage 2 versus the benchmark losses. We see that, with the concentrated simulation budget, the Stage 2 simulation significantly improves the accuracy of the loss estimates in the highly likely tail scenarios, which will in turn improve the accuracy of the CTE estimate. The subfigure



(a) Simulated losses by Stage 1 simulation ( $x$  axis) and by the benchmark simulation ( $y$  axis). (b) Simulated losses by Stage 2 simulation ( $x$  axis) and by the benchmark simulation ( $y$  axis).

Fig. 3. Illustration of the proposed two-stage simulation procedure.

Table 2

Average runtime of one repetition of each experiment, to the nearest hour.

Experiment Design	Runtime of One Repetition
Two-stage procedures	
(a1) $M^h = 1,500, N_2 = 80$	8 hours
(a2) $M^h = 1,000, N_2 = 170$	10 hours
(a3) $M^h = 500, N_2 = 620$	8 hours
(b) Standard nested, $N = 350$	32 hours

in Fig. 3b zooms into the border between tail and non-tail scenarios: In this illustration, the CTE estimate by the two-stage procedure includes 481 of the 500 true tail scenarios. Among the 19 tail scenarios that were not included in the CTE estimate, 3 were missed in Stage 1, and 16 were missed in Stage 2.

Next we examine the performance of the two-stage procedure in more detail. We repeat the following four experiments 100 times:

Experiment (a): The proposed two-stage procedure with

(a1)  $M^h = 1500$ , and  $N_2 = 80$  (the same configuration as in Figs. 3a and 3b),

(a2)  $M^h = 1000$ , and  $N_2 = 170$ , and

(a3)  $M^h = 500$ , and  $N_2 = 620$ .

Experiment (b): Standard multi-period full nested simulation with  $N = 350$  inner simulations for each outer scenario.

In Experiment (a),  $N_1$  is set equal to 2 in each case. We chose  $M^h = 500, 1000, 1500$  to test the impact of different sizes for the highly likely tail scenario set; the  $N_2$  values are then determined such that Experiments (a1), (a2), and (a3) all require similar computation time. Note that  $M^h = 500$  is the minimum size for the highly likely tail scenario set.

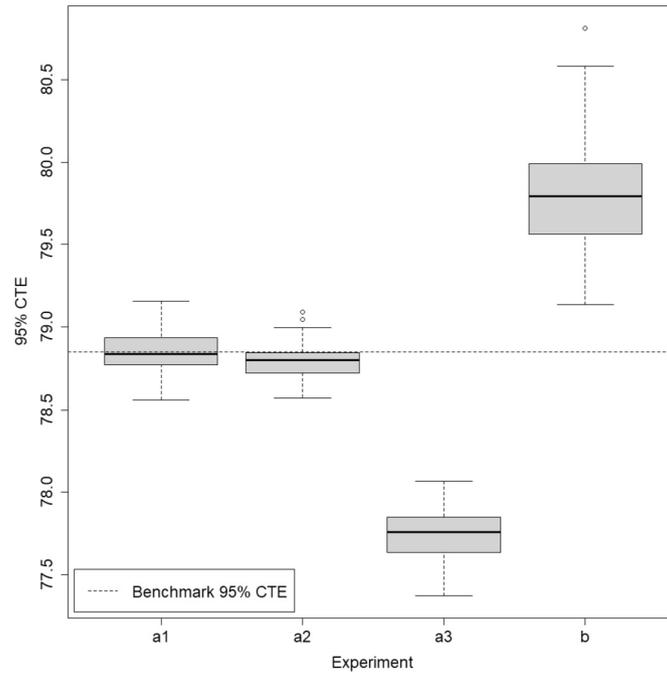
Runtime for a single repetition of the experiments shown in Table 2 is summarized in Table 2. These were conducted with 40 cores on a Dell PowerEdge R840 server with 4 Intel Xeon Gold 6230 20-core 2.1 GHz (Cascade Lake) CPU and 768 GB memory. Experiments (a1), (a2), and (a3) take around a quarter of the runtime of a standard nested simulation shown in Table 3.

The likelihood ratio computations reduce the number of inner simulations available for a given runtime. In the standard nested simulation experiment, where no likelihood computation is needed, the total number of inner sample paths in all scenarios is  $10,000 \times 350 = 3,500,000$ . In Experiment (a1), the total number of inner sample paths reduces to  $10,000 \times 2 + 1,500 \times 80 = 140,000$ , which is only 4% of that in Experiment (b), but the run time of Experiment (a1) is 25% of the run time of Experiment (b) due to the likelihood computations required.

Each of the 100 repetitions of the above experiments produces a 95%-CTE estimate. By comparing these estimates with the benchmark 95%-CTE estimate we can estimate the Root Mean Squared Errors (RMSEs), bias and standard deviation, relative to the true CTE value, for each of the four experiments.

**Table 3**  
Accuracy measures from 100 repetitions of different configurations of the two-stage procedure and the standard nested simulation.

Experiment Design	Relative RMSE	Relative Bias	Relative SD
(a) Two-stage procedures			
a1. $M^h = 1,500, N_2 = 80$	0.146%	-0.002%	0.146%
a2. $M^h = 1,000, N_2 = 170$	0.147%	-0.069%	0.130%
a3. $M^h = 500, N_2 = 620$	1.417%	-1.405%	0.182%
(b) Standard nested sim. with $N = 350$	1.263%	1.201%	0.392%



**Fig. 4.** Box-and-whisker plot of 95%-CTE estimates from 100 repetitions of three configurations of two-stage procedure using MLR estimators, and standard nested simulation.

The relative RMSE is

$$\text{Relative RMSE} = \frac{1}{CTE_\alpha} \sqrt{\frac{\sum_{i=1}^n (\widehat{CTE}_{\alpha,i} - CTE_\alpha)^2}{n}}$$

where  $n = 100$  is the number of repetitions of the experiment,  $\widehat{CTE}_{\alpha,i}$  is the estimated  $\alpha$ -CTE of the  $i$ th repeated experiment, and  $CTE_\alpha$  is the  $\alpha$ -CTE value estimated by the large scale benchmark full nested simulation, which stands in for the true  $\alpha$ -CTE for the 10,000 scenarios. The relative bias is

$$\frac{1}{CTE_\alpha} \sum_{i=1}^n \frac{(\widehat{CTE}_{\alpha,i} - CTE_\alpha)}{n},$$

and the relative standard deviation is

$$\frac{1}{CTE_\alpha} \sqrt{\frac{\sum_{i=1}^n (\widehat{CTE}_{\alpha,i} - \frac{1}{n} \sum_{i=1}^n \widehat{CTE}_{\alpha,i})^2}{n}}$$

Table 3 and Fig. 4 summarize these performance measures. We make the following observations.

- (1) Experiments (a1) and (a1) provide substantially more accurate estimates of the 95% CTE for the GMWB, with lower bias and lower variance, compared with the standard nested simulation used in Experiment (b). Recall that the run times for Experiments (a1) and (a2) are only 25% of the full nested approach used in Experiment (b).
- (2) Experiment (a3), which allows no safety margin in the highly likely tail scenario set, gives a worse estimator than (a1), (a2), or (b), in terms of the RMSE, with relatively high standard deviation and a significant negative bias. The negative bias is caused by the fact that the first step does not exactly predict the 500 true tail scenarios; a number will be misclassified and will therefore be replaced in the CTE calculation with scenarios that generate smaller losses. The second step of the inner simulation process provides more

**Table 4**  
True tail scenarios captured in 100 repetitions of two-stage nested simulation using MLR estimator of different configurations of the two-stage procedure and the standard nested simulation.

Experiment Design	# true tail scen. in $\hat{\mathcal{T}}_\alpha^{HL}$	# of repetitions $\hat{\mathcal{T}}_\alpha^{HL}$ includes all true tail scenarios	# true tail scen. in $\hat{\mathcal{T}}_{(1-\alpha)M}^{MLR}$
(a) Two-stage procedures			
a1. $N_1 = 2, M^h = 1,500, N_2 = 80$	499.6	66/100	476.9
a2. $N_1 = 2, M^h = 1,000, N_2 = 170$	498.0	17/100	478.9
a3. $N_1 = 2, M^h = 500, N_2 = 620$	438.6	0/100	438.6
(b) Standard nested sim. with $N = 350$			
	n/a	n/a	434.4

accurate loss evaluation for the scenarios selected, but the number of misclassified scenarios will inevitably decrease the estimated CTE compared with the accurate value. Experiments (a1) and (a2) use 1,500 and 1,000 Stage 1 tail scenarios, respectively, to capture the largest 500 loss values. Compared with (a1), (a2) uses more inner simulations, giving a smaller variance in the CTE estimate, but also displays a small negative bias, indicating that some of the true tail scenarios are misclassified in Stage 1, similarly to Experiment (a3), but with much less severe impact.

We examine this misclassification in more detail in Table 4, which summarizes the number of the true tail scenarios identified in Stage 1 of the two stage process, averaged over the 100 repetitions. In the second column, we show the average number of true tail scenarios that are correctly classified in Stage 1 into the highly likely tail scenario set,  $\hat{\mathcal{T}}_\alpha^{HL}$ . In the third column, we show how many times the Stage 1 process correctly classified all of the true tail scenarios into  $\hat{\mathcal{T}}_\alpha^{HL}$ . In the fourth column, we show the average number of true tail scenarios included in the final CTE calculation.

From the table we see that in Experiment (a3), none of the 100 repeated experiments correctly classifies all 500 true tail scenarios in the highly-likely tail scenario set; on average, 438.6 of the 500 true tail scenarios are correctly classified, resulting in the bias identified in Fig. 4. As each of the scenarios allocated to  $\hat{\mathcal{T}}_\alpha^{HL}$  is used in the CTE calculation, the second and fourth columns are the same.

Increasing the safety margin to 5% $M$ , as in Experiment (a2), drastically reduces the misclassification of tail scenarios. On average, 498.0 of the 500 true tail scenarios are correctly classified into the highly likely tail scenario set,  $\hat{\mathcal{T}}_\alpha^{HL}$ . In 17 of 100 repeated experiments,  $\hat{\mathcal{T}}_\alpha^{HL}$  includes all the true tail scenarios, and on average, Experiment (a2) includes 478.9 true tail scenarios in the CTE estimate.

In Experiment (a1) the number of highly-likely tail scenarios is increased to  $M^h = 1,500$ . The average number of true tail scenarios allocated to  $\hat{\mathcal{T}}_\alpha^{HL}$  is 499.6, and all of the true tail scenarios were correctly allocated in 66 of the 100 repeated experiments. However, the inner simulation step is a little less accurate than Experiment (a2), as the budget is computation budget applied more widely, resulting in a smaller number of true tail scenarios included in the final CTE calculation, on average, than Experiment (a2). As we see from Table 3, though, the slightly different misclassification between Experiments (a1) and (a2) has little effect on their RMSEs.

- (4) Experiment (b), the standard nested simulation, has high bias and high variance. This is a well documented result of using an insufficient number of inner simulations - see, for example, Gordy and Juneja (2010), Dang (2021), and Broadie et al. (2011). The explanation is that using a small number of inner simulations causes significant noise in the estimated losses for the outer scenarios, creating the high variance. In addition, because the estimated CTE is the average of the  $(1 - \alpha)M$  largest simulated losses, the losses that are underestimated because of sampling variability fall out of the CTE calculation and the losses that are over-estimated are therefore disproportionately represented in the calculation, leading to a general positive bias. In simple terms, suppose the true loss value for each of the worst 1000 scenarios is equal to 80. The 95% CTE is estimated from the average of the largest 5000 simulated losses. Now suppose that inner sampling noise will give a loss estimate for each scenario of either 70 or 90, with equal probability. Overall there is no bias, but when we average the largest 500 simulated losses to find the CTE, we get a value of 90, on average, as we disproportionately capture the high-side estimates.
- (5) Clearly, using more inner simulations produces a more accurate estimate of the loss for a given scenario. In Experiment (b) we use only 350. The estimate in Experiment (a1), uses only 82 (that is,  $N_1 + N_2$ ) original inner simulation paths for each scenario, but it also uses an additional  $1499 \times 82$  inner simulation results that are re-purposed from the other scenarios, for a total of 123,000. Similarly, including re-used simulations, Experiment (a2) uses  $1,000 \times 172 = 172,000$  inner simulation outputs for each scenario. However, the re-used simulation paths are not all equally likely, unlike the 350 paths used in Experiment (b). If the likelihood ratios are very small, meaning that the re-used simulated outputs are not very helpful across different scenarios, then the estimate might not be much better than simply using the 82 inner simulations specific to the individual scenario. A more appropriate comparison is to consider the effective sample size, which is the number of i.i.d. simulations that would have achieved the same mean squared error as the likelihood ratio estimator. This is a common diagnostic in importance sampling. In Kong (1992) and Liu (1996), under some technical assumptions, it can be shown that the effective sample size,  $\tilde{N}$ , of a likelihood ratio estimator can be estimated by

$$\tilde{N}_t^{(i)} = \frac{\left(\sum_{k=1}^M \sum_{j=1}^N w_t^{(k,j)}\right)^2}{\sum_{k=1}^M \sum_{j=1}^N \left(w_t^{(k,j)}\right)^2}$$

where  $w_t^{(k,j)}$  is the likelihood ratio used in the MLR estimator. Fig. 5 depicts the log of the average effective sample sizes for each of the experiments at different times  $t$  (in months) along the outer scenario. We see from this figure that the effective sample size for the two-stage procedure is higher than that of the standard nested simulation at all times, even in the first stage, where only two original inner simulations are generated for each of the 10,000 scenarios. In the second stage, the effective sample size is very much higher than the standard nested approach for the scenarios selected in the first stage.

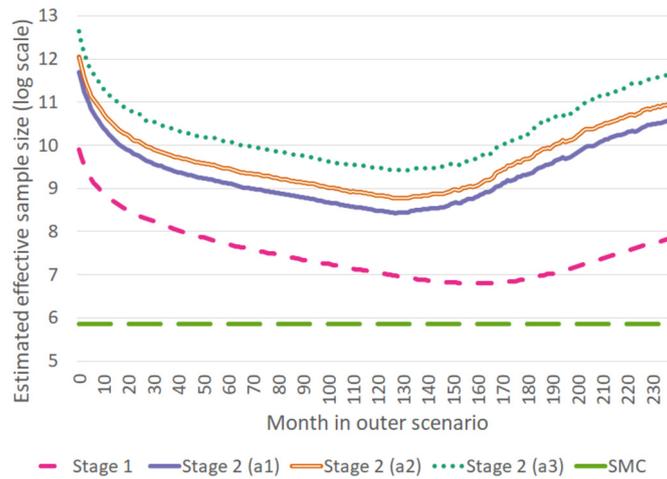


Fig. 5. Estimated effective sample size (in log scale) in one repetition of Stage 1 MLR, Stage 2 MLR, and SMC experiment.

(6) We note that Lan et al. (2010) also propose a two-stage (single-period) nested simulation procedure. We implemented their method for our sample contract and scenarios, but the results were not satisfactory. Using their screening (Stage 1) procedure to identify the 5% true tail scenarios with 95% confidence resulted in around 35% of the 10,000 scenarios being selected for Stage 2. Moreover, the pairwise student-t test used in their screening algorithm requires estimating the standard errors for  $\frac{M(M-1)}{2}$  pairs of scenarios. In our example, with  $M = 10,000$ , the screening procedure itself takes longer to run than our entire two-stage procedure, leaving no computational budget for Stage 2.

### 6. Portfolio of VA contracts

In practice, insurers' VA portfolios usually include of contracts of various types, issued to a large number of policyholders with different demographic characteristics. Interested readers may refer to Hardy (2003) for a comprehensive discussion of various types of VA contracts.

In this work, we have focused on improving the computational efficiency of a single VA contract. Nevertheless, our proposed two-stage procedure using MLR estimators can be applied in nested simulations of a portfolio of VA contracts. When considering nested simulation of a VA portfolio, each stage of our proposed procedure can be applied independently to every contract in the model, be it a contract from the original portfolio or a representative contract as proposed in, for example, Gan and Lin (2015) and Lin and Yang (2020). At the end of each stage, the highly likely tail scenario set is identified as the set of scenarios with the worst aggregate losses for the whole portfolio.

We illustrate using a numerical example how our proposed two-stage procedure using MLR estimators is applied to improve the computation efficiency of a hypothetical portfolio consisting of three VA contracts:

1. One GMWB contract as described in Section 5.
2. One GMWB contract with the same characteristic as 1, except that it has 10-years to maturity, and a monthly withdrawal rate of  $\gamma = 0.5\%$ .
3. One GMMB contract.

#### 6.1. Characteristics of a GMMB contract

The simplest form of GMMB contract pays a maturity benefit equal to the greater of the sub-account value and a fixed guarantee value. The payoff of the GMMB resembles that of a European put option.

The GMMB contract in this numerical example is modeled with the following characteristics.

- The initial fund value is  $F_0 = 1000$ , which is also the guarantee base  $G_0 = 1000$ . The guarantee base only reduces due to lapse throughout the projection.
- It has a 20-year time-to-maturity and is hedged monthly, so  $T = 240$ .
- A management fee of  $\eta_g = 0.146\%$  is deducted monthly from the fund value and of which 0.025% is treated as income per month for the contract guarantees, that is,  $\eta_n = 0.025\%$ .
- We ignore decrements from mortality.
- We assume the lapse behavior of the policyholder is dynamic in the sense that it is dependent on the moneyness of the contract. More precisely, we assume that the monthly lapse rate from  $t - 1$  to  $t$  is

$$q_{x,t-1}^l = \min \left( 1, \max \left( 0.5, 1 - 1.25 \times \left( \frac{G_{t-1}}{F_{t-1}} - 1.1 \right) \right) \right) \times q_{x,t-1}^{l-base} \tag{35}$$

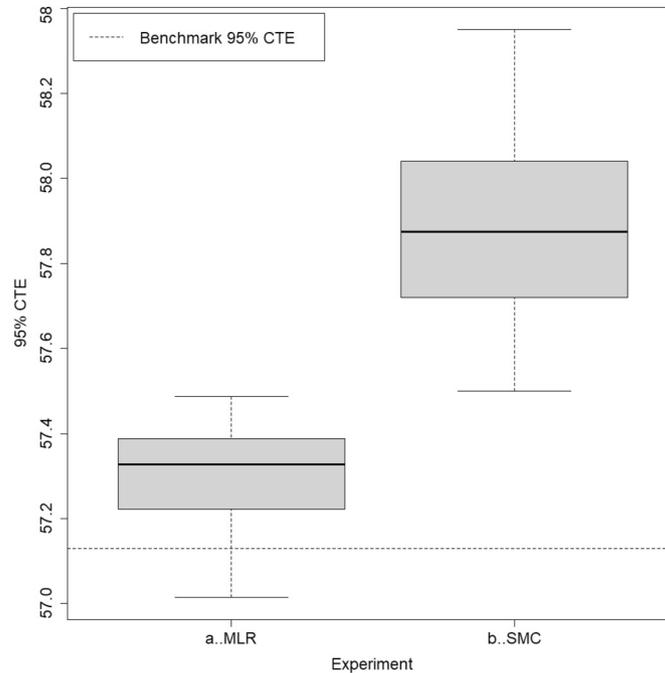
and

$$q_{x,t}^{l-base} = \begin{cases} 0.00417 & \text{if } t < 84, \\ 0.00833 & \text{if } t \geq 84. \end{cases} \tag{36}$$

The mathematical modeling of the GMMB contract is detailed in Appendix A.

**Table 5**  
Accuracy measures from 20 repetitions of the two-stage procedure for the portfolio of VA contracts.

Experiment Design	Relative RMSE	Relative Bias	Relative SD
(a) Two-stage procedure $M^h = 1,500, N_2 = 80$	0.37%	0.305%	0.21%
(b) Standard nested, $N = 350$	1.38%	1.325%	0.39%



**Fig. 6.** Box-and-whisker plot of 95%-CTE estimates of the VA portfolio from 20 repetitions of the two-stage procedure using MLR estimators, and standard nested simulation.

6.2. Numerical experiment

We considered two experiments for the 3-contract VA portfolio:

- (a) Two-stage MLR procedure with  $N_1 = 2, M^h = 1,500,$  and  $N_2 = 80$  (Same as experiment (a1) in Section 5)
- (b) Standard full nested simulation with  $N = 350, M = 10,000$  (Same as experiment (b) in Section 5)

The experiment was repeated 20 times. The relative RMSE, bias and standard deviation of the CTE estimates from each set of repeated experiments were compared to the CTE of the benchmark experiment. As in Section 5, the benchmark experiment is a standard nested simulation with  $N = 10,000$  inner simulations.

The asset model is the same as the one used in Section 5. All these experiments use the same outer scenarios. Table 5 summarizes the performance measures of this experiment. The same results are also summarized in Fig. 6.

The results in Table 5 show much higher accuracy using the two-stage procedure with MLR estimators than in standard Monte Carlo simulation, while the two-stage procedure takes roughly a quarter of the runtime of the standard nested simulation. This demonstrates the potential for applying the two-stage procedure using MLR estimators to a heterogeneous VA portfolio.

7. Concluding remarks

In this paper, we have presented a two-stage nested simulation procedure for estimating the tail risks associated with dynamic hedging of complex, path-dependent embedded options. The mixture likelihood ratio estimator is used in both stages to reuse simulation outputs and to improve the estimation accuracy. In the numerical illustrations, we apply the proposed two-stage procedure to the GMWB, and also to a heterogeneous portfolio containing GMWBs and a GMMB. Compared with a standard nested simulation, the method offers very significant improvement in the accuracy of the CTE, with much shorter run times.

There are a few potential refinements and extensions to our proposed procedure in future studies.

1. In this paper we have focused entirely on the inner simulation stage. It may be possible to develop a dynamic that trades the computation budget between inner and outer simulations.
2. The parameters  $M^h$  and  $N_1$  have been selected somewhat arbitrarily. Further research could identify more objective methods of setting these parameters.
3. The propositions in this study are mainly computational. Convergence analysis of the proposed procedure could be an area of future study.

**Declaration of competing interest**

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

**Acknowledgement**

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada, funding reference number 03754 (Hardy) and 03755 (Feng). This work was also supported by the Society of Actuaries through a Center of Actuarial Excellence Research Grant and through the Hickman Scholarship held by Ou Dang. Ou Dang has also received support from the Ontario Graduate Scholarship.

**Appendix A. Mathematical modeling of the GMMB contract detailed in Section 6.1**

At  $t = 1, \dots, T$ , the fund value  $F_t$  and guarantee value  $G_t$  of the GMMB contract evolves as follows.

$$F_t = \max \left( F_{t-1} \frac{S_t}{S_{t-1}} e^{-\eta_g} \cdot p_{x,t-1}, 0 \right) = F_{t-1} e^{R_t - \eta_g} \cdot p_{x,t-1}, \text{ where } R_t = \ln \frac{S_t}{S_{t-1}}$$

$$G_t = G_{t-1} \cdot p_{x,t-1},$$

where  $p_{x,t}$  is the probability of surviving all decrements from time  $t$  to  $t + 1$ .

At any time  $t$ , the future liability beyond time  $t$  has value of

$$v_{t+} = e^{-r(T-t)} (G_T - F_T)^+ - \sum_{s=t+1}^T e^{-r(s-t)} F_s (e^{\eta_n} - 1). \tag{37}$$

The pathwise estimator of  $\Delta_t(\mathbf{S}_t, \mathbf{F}_t, \mathbf{G}_t)$ , the number of units of the underlying stock  $S_t$  in a delta hedging program, based on a single inner simulation path, is given by

$$H_t(\tilde{\mathbf{S}}_{t+}, \tilde{\mathbf{F}}_{t+}, \tilde{\mathbf{G}}_{t+}) = -e^{-r(T-t)} \mathbb{1} \{ \tilde{G}_T > \tilde{F}_t \} \cdot \frac{d\tilde{F}_T}{dS_t} - \sum_{s=t+1}^T e^{-r(s-t)} \frac{d\tilde{F}_s}{dS_t} (e^{\eta_n} - 1), \tag{38}$$

where  $\frac{d\tilde{F}_s}{dS_t} = \frac{F_t}{S_t} e^{\sum_{u=t+1}^s (\tilde{R}_u - \eta_g)} \cdot {}_{s-t}p_{x,t}$ .

Similar to the GMWB case in Section 4.2, to conform with the evolution of the GMMB state variables, we also need to adapt the likelihood ratio method before applying it to the delta hedging of GMMB's liability. Let  $\Lambda_t := \{(\mathbf{S}_{T,i}, \mathbf{F}_{T,i}, \mathbf{G}_{T,i}) : i = 1, \dots, M, F_{t,i} > I_{t,i}\}$  be the outer scenarios that require inner simulations at time  $t$  and let  $M_t = |\Lambda_t|$ , for any target scenario  $(\mathbf{S}_{T,i}, \mathbf{F}_{T,i}, \mathbf{G}_{T,i}) \in \Lambda_t$ , the MLR estimator for the time  $t$  delta of the GMMB contract is

$$\hat{\Delta}_t^{\text{MLR}}(\mathbf{S}_{t,i}, \mathbf{F}_{t,i}, \mathbf{G}_{t,i}) = \frac{1}{M_t N} \sum_{k: (\mathbf{S}_{T,k}, \mathbf{F}_{T,k}, \mathbf{G}_{T,k}) \in \Lambda_t} \frac{p_{x,t,i}}{p_{x,t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{\text{adj}}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{F_{t,k}} \times \sum_{j=1}^N H_t(\tilde{\mathbf{S}}_{t+,kj}, \tilde{\mathbf{F}}_{t+,kj}, \tilde{\mathbf{G}}_{t+,kj}) \frac{f(\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}} | \mathbf{S}_{t,i})}{\tilde{f}_M(\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}})} \tag{39}$$

where  $\tilde{R}_{t+1,kj} = \ln \frac{\tilde{S}_{t+1,kj}}{S_{t,k}}$  and  $\tilde{R}_{t+1,kj}^{\text{adj}} = \tilde{R}_{t+1,kj} + \ln \left[ \frac{G_{t,i}}{G_{t,k}} \cdot \frac{F_{t,k}}{F_{t,i}} \right]$ .

The above MLR estimator is derived as follows. Given equation (38), we have

$$H_t^{\text{adj}}(\tilde{\mathbf{S}}_{t+,kj}, \tilde{\mathbf{F}}_{t+,kj}, \tilde{\mathbf{G}}_{t+,kj}) = -e^{-r(T-t)} \mathbb{1} \{ \tilde{G}_{T,kj}^{\text{adj}} > \tilde{F}_{T,kj}^{\text{adj}} \} \cdot \frac{d\tilde{F}_{T,kj}^{\text{adj}}}{dS_{t,i}} - \sum_{s=t+1}^T e^{-r(s-t)} \frac{d\tilde{F}_{s,kj}^{\text{adj}}}{dS_{t,i}} (e^{\eta_n} - 1) \tag{40}$$

In this equation, we have

$$\begin{aligned} & \tilde{G}_{T,kj}^{\text{adj}} > \tilde{F}_{T,kj}^{\text{adj}} \tag{41} \\ \iff & G_{t,i} \cdot {}_{T-t}p_{x,t,kj}^{\text{adj}} > F_{t,i} \cdot e^{(\tilde{R}_{u,kj}^{\text{adj}} - \eta_g) + \sum_{u=t+2}^T (\tilde{R}_{u,kj} - \eta_g)} \cdot {}_{T-t}p_{x,t,kj}^{\text{adj}} \\ \iff & G_{t,i} > F_{t,i} \cdot e^{\tilde{R}_{t+1,kj}^{\text{adj}} - \eta_g} \cdot e^{\sum_{u=t+2}^T (\tilde{R}_{u,kj} - \eta_g)} \\ \iff & G_{t,i} > F_{t,i} \cdot e^{\tilde{R}_{t+1,kj}^{\text{adj}} - \eta_g} \cdot \frac{G_{t,i}}{G_{t,k}} \cdot \frac{F_{t,k}}{F_{t,i}} \cdot e^{\sum_{u=t+2}^T (\tilde{R}_{u,kj} - \eta_g)} \text{ by } \tilde{R}_{t+1,kj}^{\text{adj}} = \tilde{R}_{t+1,kj} + \ln \left[ \frac{G_{t,i}}{G_{t,k}} \cdot \frac{F_{t,k}}{F_{t,i}} \right] \\ \iff & G_{t,k} > F_{t,k} \cdot e^{\sum_{u=t+1}^T (\tilde{R}_{u,kj} - \eta_g)} \\ \iff & G_{t,k} \cdot {}_{T-t}p_{x,t,kj} > F_{t,k} \cdot e^{\sum_{u=t+1}^T (\tilde{R}_{u,kj} - \eta_g)} \cdot {}_{T-t}p_{x,t,kj} \\ \iff & \tilde{G}_{T,kj} > \tilde{F}_{T,kj} \end{aligned}$$

**Table 6**  
Parameters for the lognormal asset model in the numerical experiments.

(Monthly rate)	Real World
Risk-free Rate: $r$	0.002
Mean: $\mu$	0.00375
Standard Deviation: $\sigma$	0.0457627

In addition, for  $s = t + 1, \dots, T$ , we have

$$\begin{aligned} \frac{d\tilde{F}_{s,kj}^{\text{adj}}}{dS_{t,i}} &= \frac{F_{t,i}}{S_{t,i}} \cdot e^{\tilde{R}_{t+1,kj}^{\text{adj}} - \eta_g} e^{\sum_{u=t+2}^s (\tilde{R}_{s,kj} - \eta_g)} \cdot p_{X,t,i} \cdot {}_{s-t-1}p_{X,t+1,kj}^{\text{adj}} \\ &= \frac{p_{X,t,i}}{p_{X,t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{\text{adj}}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}} \cdot \frac{F_{t,k}}{S_{t,k}} \cdot e^{\tilde{R}_{t+1,kj} - \eta_g} e^{\sum_{u=t+2}^s (\tilde{R}_{s,kj} - \eta_g)} \cdot p_{X,t,i} \cdot {}_{s-t}p_{X,t+1,kj} \\ &= \frac{p_{X,t,i}}{p_{X,t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{\text{adj}}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}} \cdot \frac{d\tilde{F}_{s,kj}}{dS_{t,k}} \end{aligned} \tag{42}$$

By substituting (41) and (42) into (40), we have

$$\begin{aligned} &H_t^{\text{adj}}(\tilde{\mathbf{S}}_{t+,kj}, \tilde{\mathbf{F}}_{t+,kj}, \tilde{\mathbf{G}}_{t+,kj}) \\ &= -e^{-r(T-t)} \mathbb{1} \left\{ \tilde{G}_{T,kj}^{\text{adj}} > \tilde{F}_{T,kj}^{\text{adj}} \right\} \cdot \frac{d\tilde{F}_{T,kj}^{\text{adj}}}{dS_{t,i}} - \sum_{s=t+1}^T e^{-r(s-t)} \frac{d\tilde{F}_{s,kj}^{\text{adj}}}{dS_{t,i}} (e^{\eta_n} - 1) \\ &= -e^{-r(T-t)} \mathbb{1} \left\{ \tilde{G}_{T,kj} > \tilde{F}_{T,kj} \right\} \cdot \frac{p_{X,t,i}}{p_{X,t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{\text{adj}}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}} \cdot \frac{d\tilde{F}_{s,kj}}{dS_{t,k}} \\ &\quad - \sum_{s=t+1}^T e^{-r(s-t)} \frac{p_{X,t,i}}{p_{X,t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{\text{adj}}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}} \cdot \frac{d\tilde{F}_{s,kj}}{dS_{t,k}} (e^{\eta_n} - 1) \\ &= \frac{p_{X,t,i}}{p_{X,t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{\text{adj}}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}} \cdot H_t(\tilde{\mathbf{S}}_{t+,kj}, \tilde{\mathbf{F}}_{t+,kj}, \tilde{\mathbf{G}}_{t+,kj}) \end{aligned}$$

Finally, by equation (12), we have

$$\begin{aligned} &\hat{\Delta}_t^{\text{MLR}}(\mathbf{S}_{t,i}, \mathbf{F}_{t,i}, \mathbf{G}_{t,i}) \\ &= \frac{1}{M_t N} \sum_{k: (\mathbf{S}_{T,k}, \mathbf{F}_{T,k}, \mathbf{G}_{T,k}) \in \Lambda_t} \frac{p_{X,t,i}}{p_{X,t,k}} \cdot \frac{e^{\tilde{R}_{t+1,kj}^{\text{adj}}}}{e^{\tilde{R}_{t+1,kj}}} \cdot \frac{F_{t,i}}{S_{t,i}} \times \sum_{j=1}^N H_t(\tilde{\mathbf{S}}_{t+,kj}, \tilde{\mathbf{F}}_{t+,kj}, \tilde{\mathbf{G}}_{t+,kj}) \frac{f(\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}} | \mathbf{S}_{t,i})}{\tilde{f}_M(\tilde{\mathbf{S}}_{t+,kj}^{\text{adj}})} \end{aligned}$$

**Appendix B. Numerical example using a GMMB contract**

To further demonstrate the validity and effectiveness of the proposed method, we apply the proposed two-stage procedure to estimate the CTE of a simplified Guaranteed Minimum Maturity Benefit (GMMB) contract with a log-normal asset model. It is well-known that the liability of a GMMB contract can be modeled by a put option (Hardy, 2003). With a log-normal asset model, put option value and delta can be calculated in closed-form. So, unlike the GMWB example in Section 5, inner simulation is unnecessary for GMMB experiments as it can be replaced by a closed-form calculation. In the experiments in this section, we use the closed-form calculation of delta to replace inner simulations for GMMB. The resulting benchmark CTE value is then used to estimate relative RMSE and bias, rather than resorting to the large-scale benchmark run described in the GMWB example in Section 5. As shown below, we observe similar results in the GMMB experiments as those of the GMWB experiments. In particular, compared to the standard nested simulation procedure, the proposed two-stage procedure produces more accurate CTE estimates in shorter runtime.

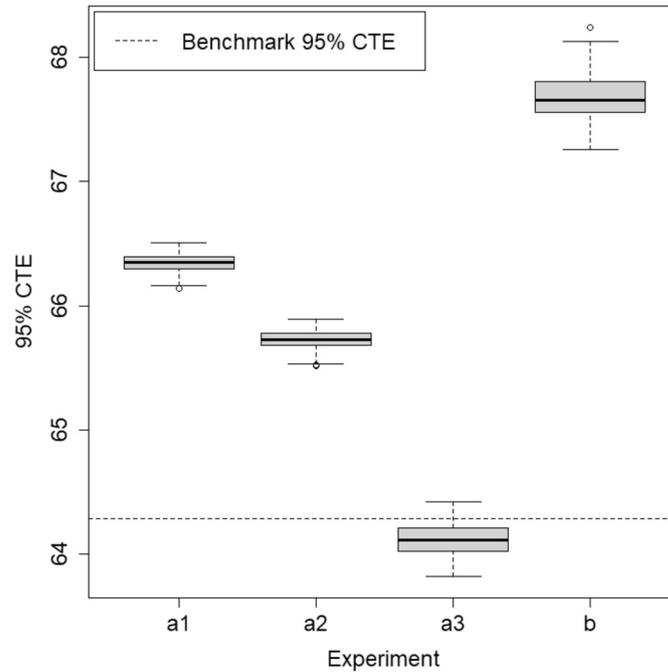
The GMMB contract we consider has the same characteristics as documented in Section 6.1 except that we ignore decrements from lapse in this experiment. We assume that the underlying stock follows a log-normal asset model, with parameters as specified in Table 6 for the real world measure.

The same four experiment designs as in Section 5, e.g., a1, a2, a3, and standard, were repeated 100 times for the GMMB contract. A fixed set of 10,000 outer scenarios were used in all experiments. The  $CTE_{95\%}$  estimate from these experiments is compared with the benchmark  $CTE_{95\%}$  estimate. To estimate the benchmark  $CTE_{95\%}$ , at each time step in the 10,000 outer scenarios, the GMMB's delta is calculated using closed-form formula. Table 7 and Fig. 7 summarize the performance measures of each set of experiment.

The results from these experiments demonstrate superior accuracy from using the two-stage compared to standard nested simulation with similar computation budget. In this GMMB experiment, experiment (a3) achieves the highest accuracy (e.g., smallest relative RMSE).

**Table 7**  
Accuracy measures from 100 repetitions of different configurations of the two-stage procedure and the standard nested simulation.

Experiment Design	Relative RMSE	Relative Bias	Relative SD
(a) Two-stage procedures			
a1. $M^h = 1,500, N_2 = 80$	3.200%	3.197%	0.130%
a2. $M^h = 1,000, N_2 = 170$	2.236%	2.233%	0.118%
a3. $M^h = 500, N_2 = 620$	0.327%	-0.260%	0.199%
(b) Standard nested sim. with $N = 350$	5.296%	5.287%	0.306%



**Fig. 7.** Box-and-whisker plot of 95%-CTE estimates from 100 repetitions of three configurations of two-stage procedure using MLR estimators, and standard nested simulation.

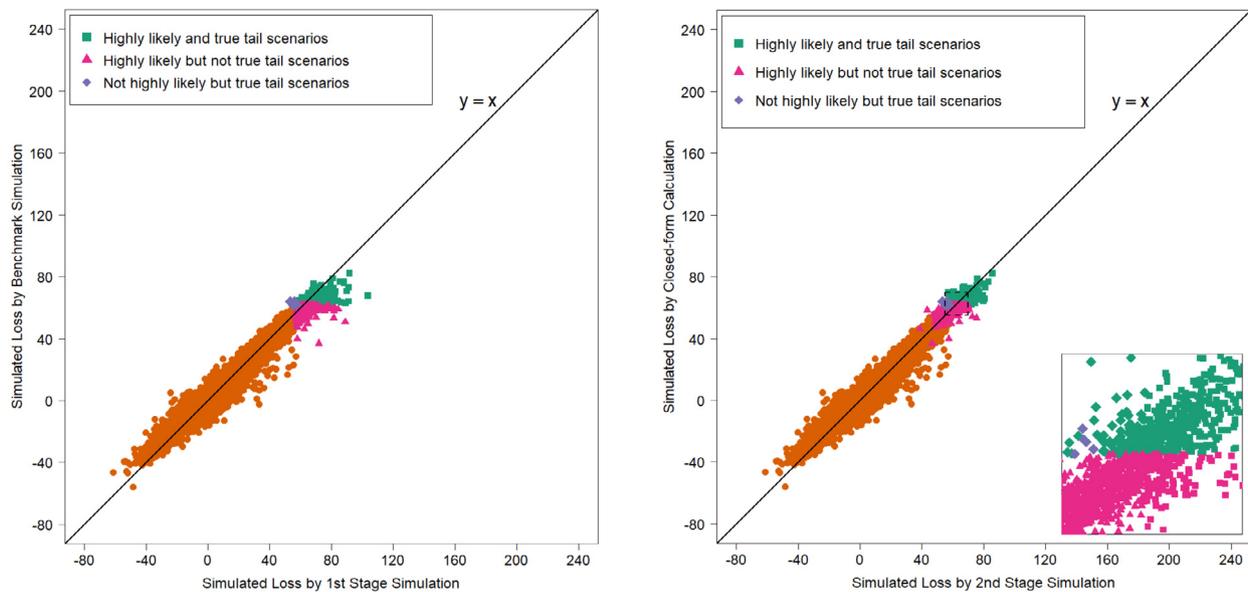
**Table 8**  
True tail scenarios captured in 100 repetitions of two-stage nested simulation using MLR estimator of different configurations of the two-stage procedure and the standard nested simulation.

Experiment Design	# true tail scen. in $\hat{\mathcal{T}}_{(1-\alpha)M}^{MLR}$
(a) Two-stage procedures	
a1. $N_1 = 2, M^h = 1,500, N_2 = 80$	385.6
a2. $N_1 = 2, M^h = 1,000, N_2 = 170$	391.4
a3. $N_1 = 2, M^h = 500, N_2 = 620$	350.7
(b) Standard nested sim. with $N = 350$	278.8

Fig. 8 shows that the MLR estimates of the losses are close to the closed-form calculation. In addition, the Stage 2 MLR estimates further improve the accuracy of the loss estimation.

As GMMB’s payoff is less complicated than GMWB’s, losses of the former in different scenarios are more clustered than the latter’s. The clustering of losses is evident in Fig. 8, which has the same x- and y-axis range as Fig. 3 but more clustered losses. This means that it is more difficult for a GMMB contract to separate the tail scenarios from non-tail scenarios than it is for a GMWB contract. Comparing Table 8 with Table 4, we see that less true tail scenarios are correctly captured in the GMMB example than in the GMWB example. The higher misclassification of tail scenarios then leads to higher relative bias, as it is evident by comparing the relative biases in Table 7 and those in Table 3.

Run times for a single repetition of the experiments shown in Table 7 are summarized in Table 9. These were conducted with 40 cores on a Dell PowerEdge R840 server with 4 Intel Xeon Gold 6230 20-core 2.1 GHz (Cascade Lake) CPU and 768 GB memory. Experiments (a1), (a2), and (a3) take around a quarter of the runtime of a standard nested simulation shown in Table 7. Compared to the GMWB experiments, the standard nested simulation takes similar run time as the complexity of its computation is driven by the number of inner simulations  $M$ , the number of outer simulations  $N$ , and the number of projection periods  $T$ . The two-stage experiments take less time than the GMWB experiments because fewer adjustments, as documented in Section 4.2 are required for this simplified GMMB contracts.



(a) Simulated losses by Stage 1 simulation ( $x$  axis) and by the closed-form calculation ( $y$  axis). (b) Simulated losses by Stage 2 simulation ( $x$  axis) and by the closed-form calculation ( $y$  axis).

Fig. 8. Illustration of the proposed two-stage simulation procedure in the GMMB example.

Table 9

Average runtime of one repetition of each experiment, to the nearest hour.

Experiment Design	Runtime of One Repetition
Two-stage procedures	
(a1) $M^h = 1,500, N_2 = 80$	3.5 hours
(a2) $M^h = 1,000, N_2 = 170$	3.6 hours
(a3) $M^h = 500, N_2 = 620$	3.6 hours
(b) Standard nested, $N = 350$	32.5 hours

In summary, Table 7 and Table 9 present the same conclusion as that for the GMWB experiment in Section 5: Compared to the standard nested simulation procedure, with well-chosen experiment design (e.g., a3) the proposed two-stage experiment can produce more accurate (e.g., 5.295% vs. 0.327% in relative RMSE) CTE estimator in shorter runtime (e.g., 32.5 hours vs. 3.6 hours).

References

Beckman, R.J., McKay, M.D., 1987. Monte Carlo estimation under different distributions using the same simulation. *Technometrics* 29 (2), 153–160.

Bollen, N.P., 1998. Valuing options in regime-switching models. *The Journal of Derivatives* 6 (1), 38–49.

Broadie, M., Du, Y., Moallemi, C.C., 2011. Efficient risk estimation via nested sequential simulation. *Management Science* 57 (6), 1172–1194.

Broadie, M., Du, Y., Moallemi, C.C., 2015. Risk estimation via regression. *Operations Research* 63 (5), 1077–1097.

Cathcart, M.J., Lok, H.Y., McNeil, A.J., Morrison, S., 2015. Calculating variable annuity liability “greeks” using Monte Carlo simulation. *ASTIN Bulletin: The Journal of the IAA* 45 (2), 239–266.

Dang, O., 2021. Efficient Nested Simulation of Tail Risk Measures for Variable Annuities.

Dang, O., Feng, M., Hardy, M.R., 2020. Efficient nested simulation for conditional tail expectation of variable annuities. *North American Actuarial Journal* 24 (2), 187–210.

Dang, O., Feng, M., Hardy, M.R., 2022. Dynamic importance allocated nested simulation for variable annuity risk measurement. *Annals of Actuarial Science* 16 (2), 319–348.

Elvira, V., Martino, L., Luengo, D., Bugallo, M.F., 2019. Generalized multiple importance sampling. *Statistical Science* 34 (1), 129–155.

Feng, M., Staum, J., 2017. Green simulation: reusing the output of repeated experiments. *ACM Transactions on Modeling and Computer Simulation* 27 (4), 1–28.

Feng, M., Staum, J., 2021. Green simulation with database Monte Carlo. *ACM Transactions on Modeling and Computer Simulation* 31 (1), 1–26.

Feng, M., Tan, Z., Zheng, J., 2020. Efficient simulation designs for valuation of large variable annuity portfolios. *North American Actuarial Journal* 24 (2), 275–289.

Feng, R., Li, P., 2022. Sample recycling method—a new approach to efficient nested Monte Carlo simulations. *Insurance. Mathematics & Economics* 105, 336–359.

Gan, G., Lin, X.S., 2015. Valuation of large variable annuity portfolios under nested simulation: a functional data approach. *Insurance. Mathematics & Economics* 62, 138–150.

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

Glasserman, P., Xu, X., 2014. Robust risk measurement and model risk. *Quantitative Finance* 14 (1), 29–58.

Gordy, M.B., Juneja, S., 2010. Nested simulation in portfolio risk measurement. *Management Science* 56 (10), 1833–1848.

Hardy, M.R., 2001. A regime-switching model of long-term stock returns. *North American Actuarial Journal* 5 (2), 41–53.

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

Hong, L.J., Juneja, S., Liu, G., 2017. Kernel smoothing for nested estimation with application to portfolio risk measurement. *Operations Research* 65 (3), 657–673.

Kleijnen, J.P., Rubinstein, R.Y., 1996. Optimization and sensitivity analysis of computer simulation models by the score function method. *European Journal of Operational Research* 88, 413–427.

Kong, A., 1992. A Note on Importance Sampling using Standardized Weights. Technical report. Department of Statistics, University of Chicago.

Lan, H., Nelson, B.L., Staum, J., 2010. A confidence interval procedure for expected shortfall risk measurement via two-level simulation. *Operations Research* 58 (5), 1481–1490.

Lee, S.H., Glynn, P.W., 2003. Computing the distribution function of a conditional expectation via Monte Carlo: discrete conditioning spaces. *ACM Transactions on Modeling and Computer Simulation* 13 (3), 238–258.

- LIMRA, 2019. LIMRA Secure Retirement Institute Annuity Sales Estimates. U.S. Individual Annuities Survey. Online (Accessed 8 December 2020).
- Lin, X.S., Yang, S., 2020. Fast and efficient nested simulation for large variable annuity portfolios: a surrogate modeling approach. *Insurance. Mathematics & Economics* 91, 85–103.
- Liu, J.S., 1996. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and Computing* 6 (2), 113–119.
- Liu, M., Staum, J., 2010. Stochastic kriging for efficient nested simulation of expected shortfall. *The Journal of Risk* 12 (3), 3.
- Maggiar, A., Wachter, A., Dolinskaya, I.S., Staum, J., 2018. A derivative-free trust-region algorithm for the optimization of functions smoothed via gaussian convolution using adaptive multiple importance sampling. *SIAM Journal on Optimization* 28 (2), 1478–1507.
- Veach, E., Guibas, L.J., 1995. Optimally combining sampling techniques for Monte Carlo rendering. In: *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*. ACM, pp. 419–428.
- Wirch, J.L., Hardy, M.R., 1999. A synthesis of risk measures for capital adequacy. *Insurance. Mathematics & Economics* 25 (3), 337–347.