

Efficient Risk Estimation via Nested Sequential Simulation*

Mark Broadie
Graduate School of Business
Columbia University
email: mnb2@columbia.edu

Yiping Du
Industrial Engineering and Operations Research
Columbia University
email: yd2166@columbia.edu

Ciamac C. Moallemi
Graduate School of Business
Columbia University
email: ciamac@gsb.columbia.edu

Initial Version: February 23, 2010
Current Revision: December 13, 2010

Abstract

We analyze the computational problem of estimating financial risk in a nested simulation. In this approach, an outer simulation is used to generate financial scenarios and an inner simulation is used to estimate future portfolio values in each scenario. We focus on one risk measure, the probability of a large loss, and we propose a new algorithm to estimate this risk. Our algorithm sequentially allocates computational effort in the inner simulation based on marginal changes in the risk estimator in each scenario. Theoretical results are given to show that the risk estimator has a faster convergence order compared to the conventional uniform inner sampling approach. Numerical results consistent with the theory are presented.

1. Introduction

The measurement and management of risk is an increasingly important function at financial institutions. A primary goal of risk measurement is to ensure that banks and other financial firms have sufficient capital reserves in relation to their holdings and investment activities. The recent failures of large and small investment and commercial banks highlight the need for better modeling and computation of financial risk measures.

Risk measurement is typically divided into two stages: scenario generation and portfolio re-valuation. Scenario generation refers to the sampling of risk factors over a given time horizon. This first (or, outer) stage is often performed with Monte Carlo simulation, especially when more realistic models with a large number of correlated risk factors are used. Portfolio re-valuation refers to the computation of the portfolio value at the risk time horizon, given a particular scenario of risk factors. Often the portfolio contains derivative securities with nonlinear payoffs that, in conjunction with more realistic financial models, require Monte Carlo simulation for this second

*This work was supported by NSF grant DMS-0914539.

(or, inner) stage. Thus, in realistic applications, the risk measurement calculation involves a two-level nested Monte Carlo simulation. Because nested Monte Carlo simulation can represent a prohibitive computational challenge, various approximation approaches are often employed. The focus of our paper is on algorithmic improvements of the *direct* nested Monte Carlo simulation approach, so that risk computation can be done on portfolios of derivative securities with more realistic multi-factor financial models.

In this paper, we consider what is perhaps the most basic risk measure, the probability that the future portfolio value falls below a pre-specified threshold, in other words, the *probability of a large loss*. When analytical formulas are available for the portfolio re-valuation step, a primary challenge of single-level Monte Carlo is to reduce the variance of the simulation risk estimator. In the nested setting, simulation is also used for the portfolio re-valuation step and additional sources of variability are introduced. The second level of simulation introduces bias into the computation, and hence both bias and variance need to be balanced and reduced to minimize the total error in the simulation risk estimate.

The problem of estimating the probability of a loss via nested simulation was first analyzed by Lee (1998) and Lee and Glynn (2003), and was subsequently considered by Gordy and Juneja (2010). These authors primarily consider and analyze *uniform* nested simulation estimators. Such estimators employ a constant number of inner samples across portfolio re-valuation calculations, thus allocating computational effort uniformly across all scenarios. They demonstrate that, asymptotically, the bias of a uniform estimator is a function of the number of inner samples used in each portfolio re-valuation, while the variance of a uniform estimator is a function of the number of outer scenarios. They characterize the asymptotically optimal uniform estimator. This estimator balances a limited computational budget between using many outer scenarios, so as to lower variance, and using many inner samples in each scenario, so as to lower bias, in a way that minimizes the overall mean squared error (MSE) among the class of uniform estimators.

This paper seeks to exploit the fact that accurate portfolio re-valuation is not equally important across all scenarios. Nested simulation can be made much more efficient by allocating computational effort *non-uniformly* across scenarios. Non-uniform estimators have been previously suggested by others in a number of contexts (e.g., Lee and Glynn, 2003; Lesnevski et al., 2004, 2007; Lan et al., 2008; Gordy and Juneja, 2008). Here, we propose and analyze a novel class of non-uniform estimators based on the idea of allocating additional effort to scenarios with a greater expected marginal change to the risk measure. Specifically, the main contributions of this paper are as follows:

1. *We propose a non-uniform nested simulation algorithm for estimating the probability of a loss.*

Our algorithm proceeds by allocating the inner stage samples for portfolio re-valuation in a sequential fashion. At each time step, it myopically selects the scenario where one additional inner stage sample will have the greatest marginal impact to the estimated loss probability. This algorithm is simple to implement and incurs minimal computational overhead.

2. *We provide an analysis that demonstrates the lower asymptotic bias of our approach.*

Given m inner stage samples in each scenario, a uniform nested estimator has an asymptotic bias of order m^{-1} . We analyze a simplified variation of our non-uniform estimator, and demonstrate that with an average of \bar{m} inner stage samples per scenario, the asymptotic bias is of order $\bar{m}^{-2+\epsilon}$, for all positive ϵ . Hence, for the same overall number of samples, the non-uniform estimator reduces bias by an order of magnitude. This theoretical analysis builds on ideas from sequential hypothesis testing, and highlights the relationship between our non-uniform estimation algorithm and classical sequential hypothesis testing.

3. *We provide an analysis that demonstrates the lower asymptotic MSE of our approach.*

Given a computational budget of k , the optimal uniform nested estimator results in an asymptotic MSE of order $k^{-2/3}$. Since non-uniform sampling provides a lower bias for the same number of inner stage samples, some of this computational savings can be used for the generation of additional outer scenarios to lower variance. We show that our non-uniform method has an asymptotic MSE of order $k^{-4/5+\epsilon}$, for all positive ϵ . Further, we demonstrate a practical implementation of our non-uniform estimator that adaptively balances bias (inner sampling) and variance (outer scenario generation).

4. *We demonstrate the practical benefits of our method via numerical experiments.*

Numerical experiments demonstrate that the performance of our non-uniform nested estimation algorithm is up to two orders of magnitude better than competing methods. Hence, we illustrate that the results achievable in practice are consistent with the gains suggested by the theory.

The rest of the paper is organized as follows. Section 1.1 contains a brief literature review. The problem setup and notation are given in Section 2. Results for uniform inner stage sampling are reviewed in Section 3. A sequential non-uniform algorithm is motivated and presented in Section 4 and a theoretical analysis is given in Section 5. Section 6 gives a practically implementable adaptive version of the sequential algorithm and numerical results are provided in Section 7. Concluding remarks are given in Section 8 and proofs are provided in the Appendix.

1.1. Literature Review

Overviews of financial risk measurement and management are given in Crouhy et al. (2000), Jorion (2006) and McNeil et al. (2006). There is a large literature on the properties of alternative risk measures (see, e.g., Artzner et al., 2000; Rockafellar and Uryasev, 2002; Föllmer and Schied, 2002). Variance reduction techniques to improve first stage sampling are given in Glasserman et al. (2000, 2002).

Most closely related to our work is that of Lee (1998) and Lee and Glynn (2003), who consider the problem of estimating the probability of a large loss and analyze nested simulation estimators and their convergence properties under uniform inner stage sampling. They consider two settings,

where the underlying scenario space is either continuous or discrete.¹ They establish that, given a total computational budget of k , the optimal uniform nested estimator results on an asymptotic MSE of order $k^{-2/3}$ in the continuous case and $k^{-1} \log k$ in the discrete case. Independently, Gordy and Juneja (2010) also consider estimating the probability of large loss in the continuous case, under a different set of assumptions. They also consider two additional risk measures (the probability of a large loss, value at risk, and expected shortfall). For each of these risk measures, they derive asymptotic bias and variance results for uniform second stage sampling. This allows them to derive the optimal allocation of effort between first and second stage sampling and derive the optimal asymptotic MSE of order $k^{-2/3}$. They also propose a jackknife procedure for reducing bias.

The idea of non-uniform nested estimation of risk measures dates back to at least the work of Lee and Glynn (2003). In the discrete case, they identify a class of non-uniform nested estimators for the probability of a large loss with asymptotic MSE of order $k^{-1} \log k$. In this setting, the non-uniform estimator achieves the same asymptotic convergence as the uniform estimator, but with a better constant. Lesnevski et al. (2004, 2007) propose a non-uniform nested estimator for a related discrete problem: they estimate the worst case expected loss across a finite set of scenarios. They are able to develop confidence intervals for their estimation procedure. Lan et al. (2007, 2008) and Liu et al. (2008) extend this work to the case of estimating expected shortfall. Contemporaneous with the present work, Gordy and Juneja (2008) suggest a broad class of non-uniform estimators for estimating the probability of a loss large, as in the present setting. Their description is rather general, however, while we provide a concrete algorithm.

Note that some of the non-uniform estimators in this prior literature have similarities to the non-uniform estimator that we propose; we discuss these in Section 4. Critically, however, none of this prior work is able to establish theoretically that a non-uniform estimator converges at a faster asymptotic *order* than is possible with uniform estimators.

There are some connections between nested simulation to estimate risk and ranking and selection (R&S) procedures which search for the best among a finite number of systems. For an overview of ranking and selection see Kim and Nelson (2005) and the book of Chen and Lee (2010). Each R&S system corresponds to an outer sample and sampling a performance measure from a system corresponds to an inner sample. Many R&S procedures rely on myopic rules to determine an allocation of inner samples (e.g., Frazier et al., 2008) and the spirit of our procedure is similar. R&S typically considers a finite and small number of systems, whereas our outer sampling draws from an infinite and often multi-dimensional domain. The R&S objective of finding the best performing system is also different than estimating a risk measure across and range of first stage outcomes.

Finally, also of interest is the work of Liu and Staum (2009); they explore an alternative approach based on stochastic kriging for estimating a risk measure. Hong and Juneja (2009) consider the benefits of kernel smoothing in risk estimation. Sun et al. (2009) consider nested simulation in the context of estimating conditional variance.

¹In this paper, we will consider only continuous scenario spaces. Note that the theory is qualitatively different in the discrete case versus the continuous case.

2. Problem Formulation: Nested Simulation

Consider the problem of measuring the risk of a portfolio of securities at some future time $t = \tau$ (the risk horizon), from the perspective of an observer at time $t = 0$. Denote the current portfolio value by X_0 . The value of the portfolio at time τ , X_τ , is, in general, a random variable and thus is not known at time 0. We assume, however, that there is a probabilistic model for the uncertainty between times 0 and τ . In particular, suppose that Ω is a set of possible future ‘scenarios’ or ‘risk factors.’ Each scenario incorporates sufficient information so as to determine all assets prices at time τ . Thus, in each scenario $\omega \in \Omega$, the portfolio has value $X_\tau(\omega)$. The *mark-to-market loss* of the portfolio at time τ in scenario ω is given by² $L(\omega) \triangleq X_0 - X_\tau(\omega)$.

A *risk measure* is a functional ρ that quantifies the risk of the random variable L by a scalar $\rho(L) \in \mathbb{R}$. Some common examples of risk measures include value-at-risk and conditional value-at-risk. In this paper, we will focus on what is perhaps the most basic risk measure, the *probability of a large loss*. That is, given a threshold $c \in \mathbb{R}$, we are interested in estimating the probability of the loss L exceeding c . Denote the resulting probability by $\alpha \triangleq \mathbb{P}(L \geq c)$.

In order to estimate the loss probability α , we face two challenges. First, typically, the space of possible scenarios Ω is quite large, if not infinite. Thus, one approach is to approximate the distribution of the loss random variable L with an empirical distribution obtained by Monte Carlo sampling. This is referred to as the *outer level* (or, first stage) of the simulation. In particular, if $\omega_1, \dots, \omega_n$ are n independent and identically distributed samples drawn according to the physical (or, real-world) distribution of ω , then we can approximate the loss probability by

$$(1) \quad \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{L(\omega_i) \geq c\}}.$$

However, even in a single scenario ω_i , it may be difficult to exactly compute the loss $L(\omega_i)$. The portfolio may contain a collection of complex, path-dependent securities with random cashflows between times τ and some final horizon T . Then, the loss $L(\omega_i)$ must be estimated via an *inner level* (or, second stage) of Monte Carlo simulation of the expected cashflows of the portfolio over the interval $[\tau, T]$. The inner simulation occurs under the risk-neutral distribution, conditioned on the scenario ω_i . If $\hat{Z}_{i,1}, \dots, \hat{Z}_{i,m}$ are m i.i.d. samples of losses generated according to this second stage of simulation, each with mean $L(\omega_i)$, then we can approximate the loss $L(\omega_i)$ in scenario ω_i by

$$(2) \quad \hat{L}_i \triangleq \frac{1}{m} \sum_{j=1}^m \hat{Z}_{i,j}.$$

The UNIFORM estimator of Algorithm 1 describes a *nested simulation* procedure that combines the estimates from the outer and inner levels of simulation in the obvious way to produce an overall

²Without loss of generality, we assume the portfolio has no intermediate cashflows before time τ , and that the riskless rate is 0.




Figure 1: Illustration of uniform sampling. The outer stage generates n financial scenarios $\omega_1, \dots, \omega_n$. Conditional on scenario ω_i , m inner stage portfolio losses $\hat{Z}_{i,1}, \dots, \hat{Z}_{i,m}$ are generated.

estimate of the loss probability. The estimator is a function of two parameters: n , the number of outer stage samples, and m , the number of inner stage samples. We say that this estimator samples *uniformly* in the sense that a constant number of inner stage samples is used for each outer stage scenario. This procedure is illustrated in Figure 1.

```

1: procedure UNIFORM( $m, n$ )
2:   for  $i \leftarrow 1$  to  $n$  do
3:     generate scenario  $\omega_i$ 
4:     conditioned on scenario  $\omega_i$ , generate i.i.d. samples  $\hat{Z}_{i,1}, \dots, \hat{Z}_{i,m}$  of portfolio losses
5:     compute an estimate of the loss in scenario  $\omega_i$ ,  $\hat{L}_i \leftarrow \frac{1}{m} \sum_{j=1}^m \hat{Z}_{i,j}$ 
6:   end for
7:   compute an estimate of the probability of a large loss,  $\hat{\alpha} \leftarrow \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{\hat{L}_i \geq c\}}$ 
8:   return  $\hat{\alpha}$ 
9: end procedure

```

Algorithm 1: Estimate the probability of a large loss using a uniform nested simulation. The parameter m is the number of inner samples per scenario. The parameter n is the number of outer scenarios.

3. Optimal Uniform Sampling

The UNIFORM estimator is a function of two parameters: n , the number of scenarios, and m , the number of inner stage samples for each scenario. This raises an obvious question: what are the best choices for the parameters m and n ? This question has been addressed in the work of Lee (1998) and Gordy and Juneja (2010). We follow the latter approach.

Denote the UNIFORM estimate of the probability of a large loss by $\hat{\alpha}_{m,n} \triangleq \text{UNIFORM}(m, n)$. The obvious objective is to choose parameters (m, n) so as to minimize the mean squared error (MSE) of the estimate $\hat{\alpha}_{m,n}$, subject to the constraint of a limited budget of computational resources. The UNIFORM estimator involves outer scenario generation and inner sampling. We will make the assumption that the computational effort of this estimator is dominated by the latter.³

³This will typically be true since the risk horizon τ is often short relative to the time horizon T of realized cashflows. In any event, the analysis in this paper can easily be extended to account for the computational effort of scenario generation.

Given parameters (m, n) , a total of mn inner samples are generated in order to compute the estimate $\hat{\alpha}_{m,n}$. Thus, given a computational work budget k on the total number of inner samples, we have the optimization problem:

$$(3) \quad \begin{aligned} & \underset{m,n}{\text{minimize}} && \mathbb{E} \left[(\hat{\alpha}_{m,n} - \alpha)^2 \right] \\ & \text{subject to} && mn \leq k, \\ & && m, n \geq 0. \end{aligned}$$

The mean squared error objective can be decomposed into variance and bias terms according to

$$(4) \quad \mathbb{E} \left[(\hat{\alpha}_{m,n} - \alpha)^2 \right] = \underbrace{\mathbb{E} \left[(\hat{\alpha}_{m,n} - \mathbb{E}[\hat{\alpha}_{m,n}])^2 \right]}_{\text{variance}} + \underbrace{\left(\mathbb{E} [\hat{\alpha}_{m,n} - \alpha] \right)^2}_{\text{bias}^2}.$$

In order to analyze the asymptotic behavior of the MSE, first consider the following technical assumption:⁴

Assumption 1. Denote by $L(\omega)$ the portfolio loss in scenario ω at time τ , and denote by \hat{L} an estimator of the form (2) for $L(\omega)$, based on the average of m i.i.d. inner stage samples. Assume that

1. The joint probability density function $p_m(\ell, \hat{\ell})$ of (L, \hat{L}) and its partial derivatives $\frac{\partial}{\partial \ell} p_m(\ell, \hat{\ell})$ and $\frac{\partial^2}{\partial \ell^2} p_m(\ell, \hat{\ell})$ exist for each m and $(\ell, \hat{\ell})$.
2. For each $m \geq 1$, there exist functions $f_{0,m}(\cdot)$, $f_{1,m}(\cdot)$, and $f_{2,m}(\cdot)$ so that

$$p_m(\ell, \hat{\ell}) \leq f_{0,m}(\hat{\ell}), \quad \left| \frac{\partial}{\partial \ell} p_m(\ell, \hat{\ell}) \right| \leq f_{1,m}(\hat{\ell}), \quad \left| \frac{\partial^2}{\partial \ell^2} p_m(\ell, \hat{\ell}) \right| \leq f_{2,m}(\hat{\ell}),$$

for all $(\ell, \hat{\ell})$. Further,

$$\sup_m \int_{-\infty}^{\infty} |\hat{\ell}|^r f_{i,m}(\hat{\ell}) d\hat{\ell} < \infty, \quad \text{for all } i = 0, 1, 2, \text{ and } 0 \leq r \leq 4.$$

Gordy and Juneja (2010) establish the following:⁵

Theorem 1. Suppose that Assumption 1 holds, and denote by $f(\cdot)$ the density of the loss variable L . As $m \rightarrow \infty$, the bias of the UNIFORM estimator asymptotically satisfies

$$\mathbb{E} [\hat{\alpha}_{m,n} - \alpha] = \frac{\theta_c}{m} + O\left(m^{-3/2}\right),$$

⁴For an alternative set of assumptions, see Lee (1998).

⁵In what follows, given arbitrary sequences $\{f_N\}$ and $\{g_N\}$, and a positive sequence $\{q_N\}$, as $N \rightarrow \infty$, we will say that $f_N = g_N + O(q_N)$ if $\limsup_{N \rightarrow \infty} |f_N - g_N|/q_N < \infty$, i.e., if the difference between f and g , is asymptotically bounded above by some constant multiple of q . Similarly, we will say that $f_N = g_N + o(q_N)$ if $\limsup_{N \rightarrow \infty} |f_N - g_N|/q_N = 0$, i.e., if the difference between f and g is asymptotically dominated by every constant multiple of q . Finally, we will say that $f_N = g_N + \Theta(q_N)$ if $0 < \liminf_{N \rightarrow \infty} |f_N - g_N|/q_N \leq \limsup_{N \rightarrow \infty} |f_N - g_N|/q_N < \infty$, i.e., if the difference between f and g is asymptotically bounded above and below by constant multiples of q .

where

$$(5) \quad \theta_c \triangleq -\Upsilon'(c), \quad \Upsilon(c) \triangleq \frac{1}{2}f(c)\mathbb{E}\left[\sigma^2(\omega) \mid L(\omega) = c\right],$$

and $\sigma^2(\omega)$ is the variance of the inner stage samples in scenario ω .

Theorem 1 directly provides an asymptotic analysis of the bias term in the MSE (4). Theorem 1 can immediately be employed to analyze the variance term, as in the following corollary:

Corollary 1. *Under the conditions of Theorem 1, as $m \rightarrow \infty$, the variance of the UNIFORM estimator satisfies*

$$\text{Var}(\hat{\alpha}_{m,n}) = \frac{\alpha(1-\alpha)}{n} + O(m^{-1}n^{-1}).$$

Proof. Note that

$$\text{Var}(\hat{\alpha}_{m,n}) = \text{Var}\left(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_{\{\hat{L}_i \geq c\}}\right) = \frac{1}{n}\text{Var}\left(\mathbb{I}_{\{\hat{L}_1 \geq c\}}\right) = \frac{\mathbb{E}[\hat{\alpha}_{m,n}](1 - \mathbb{E}[\hat{\alpha}_{m,n}])}{n},$$

where we have used the fact that the loss estimates $\{\hat{L}_i\}$ are independent and identically distributed. Applying Theorem 1,

$$\text{Var}(\hat{\alpha}_{m,n}) = \frac{\alpha(1-\alpha)}{n} + \frac{\mathbb{E}[\hat{\alpha}_{m,n} - \alpha](1 - \mathbb{E}[\hat{\alpha}_{m,n}])}{n} + \frac{\alpha\mathbb{E}[\alpha - \hat{\alpha}_{m,n}]}{n} = \frac{\alpha(1-\alpha)}{n} + O(m^{-1}n^{-1}).$$

■

Theorem 1 and Corollary 1 provide a complete asymptotic characterization of the MSE of the UNIFORM estimator. The asymptotic variance of the estimator is determined by the number of scenarios n and decays as n^{-1} , while the asymptotic bias of the estimator is determined by the number of inner stage samples per scenario m and decays as m^{-1} .

Given a computational budget of a total of k inner stage samples, a naive choice of parameters (m, n) might be to sample equally in the outer and inner stages, i.e., set $m = n = k^{1/2}$. This would result in an asymptotic bias squared of order k^{-1} and an asymptotic variance of order $k^{-1/2}$, and an overall asymptotic MSE of order $k^{-1/2}$. Since the variance is asymptotically dominating the bias squared and determining the MSE, the naive UNIFORM estimator is clearly not optimal. One could do better by using fewer inner stage samples per scenario and increasing the number of scenarios.

In order to find the optimal UNIFORM estimator, using Theorem 1 and Corollary 1, we can approximate the minimum MSE problem (3) by the optimization problem

$$\begin{aligned} & \underset{m,n}{\text{minimize}} && \frac{\alpha(1-\alpha)}{n} + \frac{\theta_c^2}{m^2} \\ & \text{subject to} && mn \leq k, \\ & && m, n \geq 0. \end{aligned}$$

This suggests optimal allocations

$$(6) \quad m^* = k^{1/3}/\beta^*, \quad n^* = \beta^* k^{2/3}, \quad \text{where } \beta^* \triangleq \left(\frac{\alpha(1-\alpha)}{2\theta_c^2} \right)^{1/3},$$

and the optimal asymptotic mean squared error

$$(7) \quad \mathbb{E} \left[(\hat{\alpha}_{m,n} - \alpha)^2 \right] = 3(\beta^*)^2 k^{-2/3} + o(k^{-2/3}).$$

The optimal allocations suggested by (6) involve, asymptotically, order $k^{2/3}$ outer stage scenarios and order $k^{1/3}$ inner stage samples per scenario. However, the optimal constant factors depend on the constant θ_c and it is not clear how to effectively estimate θ_c a priori. As we will see in Section 7, the choice of these constant factors is critical to the practical performance of a uniform estimator.

Finally, it is instructive to compare the rate of convergence of the optimal UNIFORM estimator in a two-level nested Monte Carlo simulation to that of an estimator of the probability of a large loss in a *single-level* Monte Carlo simulation. In the latter case, scenarios $\omega_1, \dots, \omega_n$ are generated. It is assumed that in each scenario ω_i , the loss $L(\omega_i)$ can be exactly computed, and the probability is estimated via (1). Note that the estimator (1) is unbiased, and has a variance proportional to n^{-1} . In a single-level simulation, then, the amount of work is proportional to n , while the MSE of the estimator decays proportional to n^{-1} . In a two-level simulation, however, as shown above, the amount of work is proportional to k , while the MSE decays at best at a rate of $k^{-2/3}$. This slower rate of decay is due to the bias introduced by the inner level of simulation.

4. Sequential Sampling

The UNIFORM estimator described in Sections 2 and 3 employs a constant number of inner stage samples for each outer stage sample. It is intuitively clear to see that this may not be an efficient strategy. As an illustrative example, consider the situation depicted in Figure 2. Here, we wish to estimate the loss probability associated with the shaded region. There are two outer stage scenarios ω_1 and ω_2 , associated with the portfolio losses $L(\omega_1)$ and $L(\omega_2)$ respectively. These true losses are approximated, in each scenario, by the estimated losses \hat{L}_1 and \hat{L}_2 .

Suppose that, under a uniform nested simulation, the portfolio losses estimated in each scenario are distributed according to the dashed probability distributions. Then, it is clear that it would be advantageous to employ fewer inner stage samples at scenario ω_1 , and more inner stage samples at scenario ω_2 . This is because the loss probability estimate $\hat{\alpha}$ is calculated according to

$$(8) \quad \hat{\alpha} \triangleq \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{\hat{L}_i \geq c\}}.$$

Thus, only the ordinal position of the estimates \hat{L}_1 and \hat{L}_2 relative to the loss threshold c is relevant. Given the uncertainty in the estimate \hat{L}_1 , it is fairly certain that $L(\omega_1) < c$, and, indeed, this could

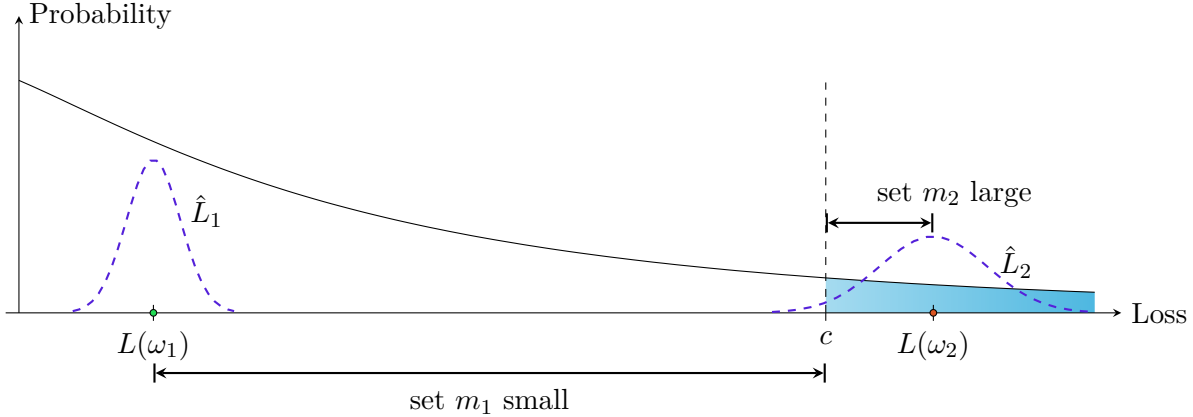


Figure 2: An illustration of the benefits of non-uniform sampling. The uncertainty in the loss \hat{L}_1 estimated in scenario ω_1 is unlikely to impact the overall probability of large loss estimate, hence the number of inner samples m_1 in this scenario can be set small. In scenario ω_2 , however a large number of inner samples m_2 should be used.

likely be inferred using *fewer* inner samples in scenario ω_1 . Given the uncertainty in the estimate \hat{L}_2 , the fact that $L(\omega_2) \geq c$, on the other hand, is much less certain. Without *more* inner samples in this scenario, there may be significant risk of misclassifying $L(\omega_2)$. These observations suggest that a *non-uniform* sampling strategy may be superior: the number of inner samples m_1 employed at scenario ω_1 should be less than the number of inner samples m_2 employed at scenario ω_2 .

The discussion above suggests that in a scenario ω with a loss $L(\omega)$ that is much greater than c or much less than c , few inner samples are necessary. If the loss $L(\omega)$ is close to c , however, many inner samples are necessary. Unfortunately, a priori, it is not clear how to do this. It is impossible to know the value of $L(\omega)$ — this is exactly what we seek to estimate via the inner Monte Carlo simulation.

We propose a procedure that simultaneously maintains estimates of the loss in each scenario, while *sequentially* attempting to allocate additional inner samples across the outer scenarios. We will first motivate our algorithm with an informal justification, and then give a precise description. In particular, suppose that there are n scenarios $\omega_1, \dots, \omega_n$. For each scenario ω_i , suppose that m_i inner samples $\hat{Z}_{i,1}, \dots, \hat{Z}_{i,m_i}$ have been made, resulting in the loss estimate $\hat{L}_i \triangleq \frac{1}{m_i} \sum_{j=1}^{m_i} \hat{Z}_{i,j}$. This results in an overall probability of a large loss estimate $\hat{\alpha}$ given by (8).

Without loss of generality, assume that $\hat{L}_i \geq c$. Suppose we wish to perform one additional inner stage sample. If we were to perform the additional sample in scenario ω_i , this would result in a new loss estimate given by

$$\hat{L}'_i \triangleq \frac{1}{m_i + 1} \sum_{j=1}^{m_i+1} \hat{Z}_{i,j} = \frac{1}{m_i + 1} \hat{Z}_{i,m_i+1} + \frac{m_i}{m_i + 1} \hat{L}_i.$$

The additional sample will only impact the estimate $\hat{\alpha}$ if the \hat{L}_i is on the opposite side of the threshold level c than \hat{L}'_i , i.e., if $\hat{L}'_i < c$. This is illustrated in Figure 3. In order to myopically

maximize the impact of the single additional sample, we will seek to choose the scenario ω_i which maximizes the probability of such a sign change. Suppose that the additional sample \hat{Z}_{i,m_i+1} has variance $\sigma_i^2 \triangleq \sigma^2(\omega_i)$. Observe that

$$(9) \quad \begin{aligned} \mathbb{P}(\hat{L}'_i < c) &= \mathbb{P}\left(\hat{Z}_{i,m_i+1} - L(\omega_i) < -m_i(\hat{L}_i - c) - (L(\omega_i) - c)\right) \\ &\approx \mathbb{P}\left(\hat{Z}_{i,m_i+1} - L(\omega_i) < -m_i|\hat{L}_i - c\right) \leq \left(1 + \frac{m_i^2}{\sigma_i^2}|\hat{L}_i - c|^2\right)^{-1}. \end{aligned}$$

Here, the approximation follows from the assumption that $m_i \gg 1$, so that $-m_i(\hat{L}_i - c) - (L(\omega_i) - c) \approx -m_i|\hat{L}_i - c|$. The inequality follows from the one-sided Chebyshev inequality. By analogous consideration of the symmetric case (where $\hat{L}_i < c$), a myopic allocation rule that seeks to maximize the probability of a sign change estimated via the Chebyshev bound⁶ (9) will choose to add the additional inner sample in scenario ω_{i^*} where

$$(10) \quad i^* \in \operatorname{argmin}_i \frac{m_i}{\sigma_i} |\hat{L}_i - c|.$$

An alternative justification for the myopic rule (10) arises if the additional sample \hat{Z}_{i,m_i+1} is drawn from a *location-scale* family of distributions, e.g., if \hat{Z}_{i,m_i+1} is normally distributed. Such a distribution is specified by a mean $L(\omega_i)$ and a variance σ_i^2 , so that

$$\mathbb{P}\left(\hat{Z}_{i,m_i+1} < z\right) = G\left(\frac{z - L(\omega_i)}{\sigma_i}\right),$$

where G is an increasing function. In this case,

$$(11) \quad \mathbb{P}(\hat{L}'_i < c) \approx \mathbb{P}\left(\hat{Z}_{i,m_i+1} - L(\omega_i) < -m_i|\hat{L}_i - c\right) = G\left(-\frac{m_i}{\sigma_i}|\hat{L}_i - c\right).$$

Maximizing the probability of a sign change according to (11) also results in the myopic rule (10).

We call the quantity minimized in (10), $(m_i/\sigma_i)|\hat{L}_i - c|$, the *error margin* associated with the scenario ω_i . The allocation rule (10), which picks a scenario by greedily minimizing the error margin, makes intuitive sense qualitatively. It encourages additional inner samples at scenarios which are close to the loss boundary (i.e., $|\hat{L}_i - c|$ is small), scenarios with few inner samples (i.e., m_i is small), or scenarios with significant variability in the portfolio losses (i.e., σ_i is large). The SEQUENTIAL estimator of Algorithm 2 employs the allocation rule (10). This estimator takes a triple (m^0, \bar{m}, n) of input parameters. Here, n is the desired number of outer stage scenarios, m^0 is the *initial* number of inner stage samples per scenario, and \bar{m} is the desired *average* number of inner stages samples per scenario at the conclusion of the algorithm. The algorithm proceeds as follows: first, n scenarios are generated, and, for each scenario, m^0 inner stage samples are performed. The remaining $\bar{m}n - m^0n$ inner stage samples are allocated one at a time in a sequential fashion myopically, as in (10).

⁶We thank an anonymous reviewer for suggesting this motivation.




Figure 3: An additional inner sample in scenario ω_i will only change the overall probability of loss estimate if \hat{L}_i changes sign.

```

1: procedure SEQUENTIAL( $m^0, \bar{m}, n$ )
2:   for  $i \leftarrow 1$  to  $n$  do
3:     generate scenario  $\omega_i$ 
4:     conditioned on scenario  $\omega_i$ , generate i.i.d. samples  $\hat{Z}_{i,1}, \dots, \hat{Z}_{i,m^0}$  of portfolio losses
5:      $m_i \leftarrow m^0$ 
6:   end for
7:   while  $\sum_{i=1}^n m_i < \bar{m}n$  do
8:     set  $i^* \in \operatorname{argmin}_i m_i |\hat{L}_i - c| / \sigma_i$ , where, for each  $1 \leq i \leq n$ ,  $\hat{L}_i$  is the current estimate of
     the loss in scenario  $\omega_i$ ,  $\hat{L}_i \leftarrow \frac{1}{m_i} \sum_{j=1}^{m_i} \hat{Z}_{i,j}$ , and  $\sigma_i$  is the standard deviation of the distribution
     of losses in scenario  $\omega_i$ 
9:     generate one additional portfolio loss sample  $\hat{Z}_{i^*, m_{i^*} + 1}$  in scenario  $\omega_{i^*}$ 
10:     $m_{i^*} \leftarrow m_{i^*} + 1$ 
11:  end while
12:  compute an estimate of the probability of a large loss,  $\hat{\alpha} \leftarrow \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{\hat{L}_i \geq c\}}$ 
13:  return  $\hat{\alpha}$ 
14: end procedure

```

Algorithm 2: Estimate the probability of a large loss using a sequential non-uniform nested simulation. The parameter m^0 is the initial number of inner samples per scenario. The parameter \bar{m} is the average number of inner samples per scenario at the conclusion of the simulation. The parameter n is the number of scenarios.

Note that the SEQUENTIAL estimator requires access to the conditional standard deviation σ_i^2 of losses in each scenario ω_i , in order to compute the error margin. These are not required for the UNIFORM estimator and, moreover, are typically not known in practice. However, these conditional standard deviations can be estimated in an online fashion over the course of the estimation algorithm; we discuss such variations in Section 7.4.

Further, the SEQUENTIAL estimator requires additional computational overhead beyond that of the UNIFORM estimator. However, this is minimal: the only additional requirement is to track scenarios in order of error margin. This can be accomplished efficiently via a priority queue data structure (see, e.g., Cormen et al., 2002). With a priority queue, determining the scenario with minimum error margin (line 8 in Algorithm 2) can be accomplished in constant time (i.e., in an amount of time independent of m and n). Once a new inner sample is generated for a scenario (lines 9–10 in Algorithm 2), order $\log n$ time would be required to update the priority queue data structure. In practice, this is not significant.

The SEQUENTIAL estimator also requires more memory than the UNIFORM estimator. In particular, the UNIFORM estimator can be implemented in a way where scenarios are processed one-at-a-time and never need to be simultaneously stored in memory. Such an implementation would have a constant memory requirement (i.e., independent of m and n). For the SEQUENTIAL estimator, each of the n outer scenarios must be stored in memory over the course of the algorithm, hence the memory requirement is of order n . In practice, even given a very large number of scenarios (e.g., millions), each of very high dimension (e.g., thousands), this memory requirement is well within the reach of commodity hardware. Each inner sample may require simulating multiple steps over a long time horizon, but the memory requirement is minimal since all intermediate computations are discarded and only the inner sample loss is recorded.⁷

The SEQUENTIAL estimator has some similarities to non-uniform estimators that have been proposed in the literature. Lee and Glynn (2003) suggest a non-uniform nested estimator in the case where the scenario space is discrete. They choose the number of inner samples m_i in each scenario ω_i so as to optimize certain large deviation asymptotics. Using a Gaussian approximation as a heuristic, this results in the allocation

$$(12) \quad m_i \propto \frac{\sigma_i^2}{(L(\omega_i) - c)^2}.$$

Since the loss $L(\omega_i)$ in scenario ω_i is unknown, Lee and Glynn (2003) propose a two-pass algorithm: in the first pass, a small number of inner samples are generated in each scenario and are used to compute inner sample allocations in a second “production run.”

Our SEQUENTIAL estimator differs from (12) in several fundamental ways: first, the allocation (12) is loosely analogous to minimizing the square of the error margin, as opposed to the error margin itself. Second, the allocation (12) is accomplished with multiple passes, while our estimator

⁷The non-uniform THRESHOLD estimator that will be discussed in Section 5.1 does not require any additional computational or memory overhead beyond that of the standard UNIFORM estimator.

is fully sequential. Indeed, in Section 5, tools from sequential analysis will prove fundamental in the theoretical analysis of our estimators. Finally, and most importantly, in the setting of Lee and Glynn (2003), non-uniform sampling does not provide a qualitatively different rate of convergence than uniform sampling. Given a total computational budget of order k , both the uniform and non-uniform methods achieve an asymptotic MSE of order $k^{-1} \log k$, albeit with different constants. As we shall see in Section 5, we will be able to establish theoretically that a non-uniform estimator converges at a faster asymptotic *order* than is possible with uniform estimators.

Gordy and Juneja (2008) suggest a general class of multi-pass “dynamic allocation” schemes for non-uniform nested estimation. Such schemes would, for example, divide the simulation into a sequence of J phases, where in the j th phase inner samples would only be allocated to scenarios ω_i if $\hat{L}_i \geq c - \epsilon_j$. Here, $\epsilon_1 > \epsilon_2 > \dots > \epsilon_J$ are a sequence of thresholds. Gordy and Juneja (2008) provide some numerical evidence that such schemes may provide a significant improvement over uniform estimators, but the choice of specific parameters of the algorithm (e.g., the number of phases J or the thresholds $\{\epsilon_j\}$) is left as a future direction of research.

5. Analysis

In Section 4, we introduced the non-uniform SEQUENTIAL estimator and motivated this algorithm via an informal discussion. In this section, we will provide an analysis of non-uniform estimation. We begin in Section 5.1 by introducing a simplified variation of the SEQUENTIAL estimator. This simplified estimator preserves the myopic and non-uniform behavior of the SEQUENTIAL estimator, but is more amenable to analysis. Moreover, the simplified estimator is reminiscent of a compound sequential hypothesis test, and highlights connections to the classical field of sequential analysis. In Section 5.2, we provide an asymptotic analysis of the bias and variance of simplified non-uniform estimator. Finally, in Section 5.3, we discuss optimal parameter choices for the simplified non-uniform estimator. We demonstrate that this estimator has asymptotic MSE of order $k^{-4/5+\epsilon}$, for all positive ϵ , as a function of the computational budget k . This can be compared to the asymptotic MSE of order $k^{-2/3}$ of the optimal uniform estimator.

5.1. A Simplified Non-Uniform Estimator

Analysis of the SEQUENTIAL estimator described in Section 4 presents a number of challenges. Foremost among these is the fact that, over the course of the nested simulation of the SEQUENTIAL estimator, the loss estimates $\hat{L}_1, \dots, \hat{L}_n$ are dependent random variables. This dependence is induced by the myopic selection rule (10), which, at each point in time, simultaneously depends upon *all* of the loss estimates. In order to make the analysis tractable, we will consider a modification of the SEQUENTIAL estimator which results in *independent* loss estimates, while maintaining the spirit of myopic non-uniform sampling.

In particular, recall that the SEQUENTIAL estimator takes as input a parameter \bar{m} , specifying the desired average number of inner samples in each scenario, and a parameter n , specifying the

desired number of scenarios. Over the course of the algorithm, $\bar{m}n$ total inner stage samples will be generated. These samples are allocated in a sequential fashion so as to myopically minimize the error margin $(m_i/\sigma_i)|\hat{L}_i - c|$, uniformly over $1 \leq i \leq n$.

If we imagine the algorithm to be in a state where a significant number of inner samples have been generated, i.e., $m_i \gg 1$ for each i , then one would expect the error margins to be roughly constant — if not, more inner samples would have been generated for the scenarios with lower error margins. One could achieve a similar effect by fixing a *threshold* $\gamma > 0$, and continuing to add inner stage samples to each scenario ω_i until the error margin exceeds γ , i.e.,

$$(13) \quad \frac{m_i}{\sigma_i} |\hat{L}_i - c| \geq \gamma.$$

This is precisely what is done by the THRESHOLD estimator of Algorithm 3.

```

1: procedure THRESHOLD( $\gamma, n$ )
2:   for  $i \leftarrow 1$  to  $n$  do
3:     generate scenario  $\omega_i$ 
4:     set  $\sigma_i$  to be the standard deviation of the distribution of the losses in scenario  $\omega_i$ 
5:      $m_i \leftarrow 0$ 
6:     repeat
7:       generate one additional portfolio loss sample  $\hat{Z}_{i,m_i+1}$  in scenario  $\omega_i$ 
8:        $m_i \leftarrow m_i + 1$ 
9:       compute an estimate of the loss in scenario  $\omega_i$ ,  $\hat{L}_i \leftarrow \frac{1}{m_i} \sum_{j=1}^{m_i} \hat{Z}_{i,j}$ 
10:    until  $\frac{m_i}{\sigma_i} |\hat{L}_i - c| \geq \gamma$ 
11:   end for
12:   compute an estimate of the probability of a large loss,  $\hat{\alpha} \leftarrow \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{\hat{L}_i \geq c\}}$ 
13:   return  $\hat{\alpha}$ 
14: end procedure

```

Algorithm 3: Estimate the probability of a large loss using a threshold-based non-uniform nested simulation. The parameter γ is the error margin threshold. The parameter n is the number of scenarios.

At a high level, the SEQUENTIAL and THRESHOLD estimators are quite similar. Both seek to non-uniformly allocate inner stage samples based on minimization of the error margin. However, they are parameterized differently. The SEQUENTIAL estimator takes as an input the parameter \bar{m} , which is the mean number of inner stage samples. On the other hand, the THRESHOLD estimator takes as input the parameter γ , which is the threshold for the error margin. As argued earlier, for large values of \bar{m} and γ , these two algorithms yield similar results. Further, we will see numerical evidence for this in Section 7.

From a practical perspective, the SEQUENTIAL estimator is more natural. In particular, if all other parameters are fixed, it is easy to choose a value for \bar{m} . This parameter explicitly specifies the total number of inner stage samples to be generated by $\bar{m}n$, and therefore determines the running time of the algorithm. Thus, we can choose \bar{m} based on the available running time. In the THRESHOLD estimator, the parameter γ *implicitly* specifies the total number of inner stage samples

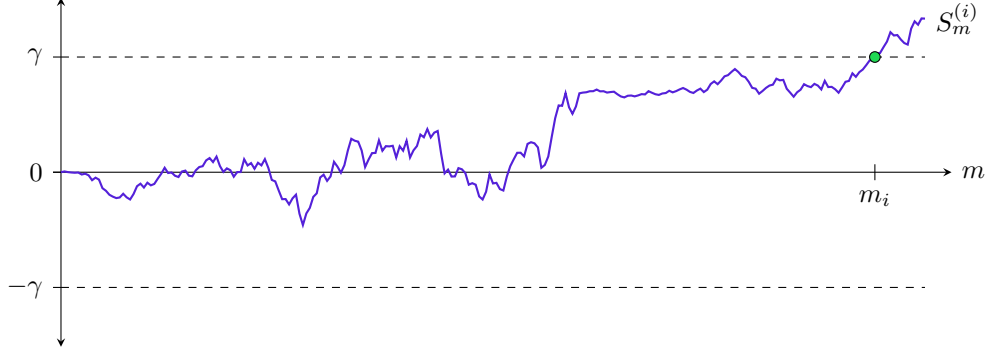


Figure 4: An illustration of the threshold estimator. Given a scenario ω_i , the estimator generates inner stage samples until the partial sum $S_m^{(i)}$ crosses barriers at γ or $-\gamma$. If the exit occurs through the upper barrier at γ , as illustrated, the scenario is declared to be a loss exceeding c . If the exit occurs through the lower barrier at $-\gamma$, the scenario is declared *not* to be a loss exceeding c .

to be generated, and hence indirectly determines the running time. It is not clear, however, how to make choice of γ a priori that ensure a certain running time, for example.

From a theoretical perspective, however, the THRESHOLD estimator proves much more amenable to analysis. The main reason is that, at any point during the execution of the algorithm, the loss estimates $\hat{L}_1, \dots, \hat{L}_n$ are *independent and identically distributed* random variables. This i.i.d. structure will prove crucial in the analysis of Section 5.2, as it allows the analysis of the overall algorithm via the analysis of a single outer stage scenario.

Moreover, the THRESHOLD estimator has another interesting interpretation. Given a threshold γ , consider a scenario ω_i with inner loss samples $\hat{Z}_{i,1}, \hat{Z}_{i,2}, \dots$. Examining (13), the algorithm will generate m_i inner stage samples in this scenario, with

$$(14) \quad m_i = \inf \left\{ m > 0 : \left| S_m^{(i)} \right| \geq \gamma \right\},$$

where, for $m \geq 0$, the partial sum is defined by

$$(15) \quad S_m^{(i)} \triangleq \sum_{j=1}^m \frac{1}{\sigma_i} \left(\hat{Z}_{i,j} - c \right).$$

Note that $\{S_m^{(i)}, m \geq 0\}$ is a random walk with unit variance increments. Then, the number of samples m_i is determined by the first exit time of the random walk from the interval $(-\gamma, \gamma)$. This is illustrated in Figure 4. If the exit occurs through the upper barrier at γ , then $\hat{L}_i > c$ and the scenario is declared to be a loss exceeding c . If the exist occurs through the lower barrier at $-\gamma$, then $\hat{L}_i < c$ and the scenario is declared *not* to be a loss exceeding c .

The interpretation of the threshold policy in terms of the first exit of a random walk is reminiscent of sequential hypothesis testing (see, e.g., Siegmund, 1985). Indeed, for each scenario ω_i , the threshold estimator is defining a sequential compound hypothesis test of whether the i.i.d. unit variance random variables $\{(Z_{i,j} - c)/\sigma_i\}$ have positive or negative mean. As we show next,

techniques from sequential analysis will prove helpful in theoretical analysis of our algorithm.

5.2. Asymptotic Analysis

Define $\tilde{\alpha}_{\gamma,n}$ to be the THRESHOLD estimate, i.e., $\tilde{\alpha}_{\gamma,n} \triangleq \text{THRESHOLD}(\gamma, n)$. As in Section 3, we will analyze the accuracy of this estimator by decomposing the mean squared error into bias and variance terms. We begin with an assumption:

Assumption 2. *Assume that:*

1. *Conditional on an outer stage scenario $\omega_i \in \Omega$, the inner stage samples $\hat{Z}_{i,1}, \hat{Z}_{i,2}, \dots$ are i.i.d. normal random variables. Denote the standard deviation of these samples by $\sigma(\omega_i)$.*
2. *Given a scenario $\omega \in \Omega$, define the normalized excess loss $\mu(\omega) \triangleq (L(\omega) - c)/\sigma(\omega)$. Then, the probability density function p of μ ,*

$$p(u) \triangleq \frac{d}{du} \mathbf{P}(\mu \leq u),$$

exists and is continuously differentiable in a neighborhood of 0.

The second condition of Assumption 2 is a technical condition that is reminiscent of the first condition of Assumption 1. The first condition is motivated by the random walk interpretation of Section 5.1. In particular, consider the random walk formed by the partial sums $\{S_m^{(i)}, m \geq 0\}$ from (15). By the functional central limit theorem, under a proper scaling, this process converges to a Brownian motion, i.e., a random walk with normal increments. The first condition makes the assumption that the unscaled random walk also has normal increments.

We are interested in the accuracy of the THRESHOLD estimator in the asymptotic regime where the resulting estimate converges to the true value, i.e., as $n \rightarrow \infty$ (many outer stage scenarios) and $\gamma \rightarrow \infty$ (many inner stage samples). Our first result is the following theorem, which characterizes the asymptotic bias of this estimator:

Theorem 2. *Under Assumption 2, as $\gamma \rightarrow \infty$, the asymptotic bias of the THRESHOLD estimator satisfies $\mathbf{E}[\tilde{\alpha}_{\gamma,n} - \alpha] = O(\gamma^{-2})$.*

The proof of Theorem 2 is provided in Appendix A. It relies on the random walk interpretation of Section 5.1 as well as techniques from sequential analysis. Specifically, exponential martingales are used in combination with the optional stopping theorem.

The following is an immediate corollary of Theorem 2, and provides an asymptotic expression for the variance of the simplified sequential estimator:

Corollary 2. *Under the conditions of Theorem 2, as $\gamma \rightarrow \infty$, the variance of the THRESHOLD estimator satisfies*

$$\text{Var}(\tilde{\alpha}_{\gamma,n}) = \frac{\alpha(1-\alpha)}{n} + O(\gamma^{-2}n^{-1}).$$

Proof. Note that

$$\text{Var}(\tilde{\alpha}_{\gamma,n}) = \text{Var}\left(\frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{\hat{L}_i \geq c\}}\right) = \frac{1}{n} \text{Var}\left(\mathbb{I}_{\{\hat{L}_1 \geq c\}}\right) = \frac{\mathbb{E}[\tilde{\alpha}_{\gamma,n}](1 - \mathbb{E}[\tilde{\alpha}_{\gamma,n}])}{n},$$

where we have used the fact that the loss estimates $\{\hat{L}_i\}$ are independent and identically distributed. Applying Theorem 2,

$$\text{Var}(\tilde{\alpha}_{\gamma,n}) = \frac{\alpha(1-\alpha)}{n} + \frac{\mathbb{E}[\tilde{\alpha}_{\gamma,n} - \alpha](1 - \mathbb{E}[\tilde{\alpha}_{\gamma,n}])}{n} + \frac{\alpha \mathbb{E}[\alpha - \tilde{\alpha}_{\gamma,n}]}{n} = \frac{\alpha(1-\alpha)}{n} + O(\gamma^{-2}n^{-1}).$$

■

The total run-time of the THRESHOLD estimator is proportional to the total number of inner stage samples generated. Note, however, by the nature of the algorithm, the number of inner samples is stochastic. Hence, define $\bar{m}(\gamma)$ to be the expected number of inner stage samples at a single outer stage scenario, given parameter γ . That is,

$$(16) \quad \bar{m}(\gamma) \triangleq \mathbb{E}\left[\inf\left\{m > 0 : \frac{m}{\sigma(\omega)} \left|\hat{L}(\omega) - c\right| \geq \gamma\right\}\right].$$

Here, the expectation is over the scenario ω and the corresponding loss estimate $\hat{L}(\omega)$. Then, given parameters (γ, n) , the THRESHOLD estimator has expected run-time proportional to $\bar{m}(\gamma)n$. The following theorems, whose proof is given in Appendix A, characterizes the rate of growth of this run-time, as a function of γ :

Theorem 3. *Under Assumption 2, as $\gamma \rightarrow \infty$, the expected number of inner stages samples in each scenario under the THRESHOLD estimator satisfies $\bar{m}(\gamma) = O(\gamma \log \gamma)$.*

Note that Theorem 3 is intuitive given the first exit time interpretation of Figure 4. In particular, for large values of γ , the amount of time required for a random walk starting at the origin with drift $\mu \neq 0$ to exit the interval $(-\gamma, \gamma)$ is approximately $\gamma/|\mu|$. If the random walk has zero drift, the exit time is approximately γ^2 . In our case, the expected number of samples $\bar{m}(\gamma)$ is averaged over various possibilities of drift given by $\mu(\omega) \triangleq (L(\omega) - c)/\sigma(\omega)$. The probability of this drift being exactly zero is zero, by the second condition of Assumption 2. However, arbitrarily small drifts are possible and thus $\bar{m}(\gamma)$ is slightly larger than $O(\gamma)$.

While Theorem 3 provides a $O(\gamma \log \gamma)$ bound on the *expected* number of inner stage samples per scenario, it might be the case that the *realized* number of inner stage samples per scenario is larger. The following theorem guarantees that, so long as the number of scenarios n is sufficiently large, a $O(\gamma \log \gamma)$ bound continues to hold on the number of realized samples per scenario with high probability. The proof can be found in Appendix A.

Theorem 4. *Under Assumption 2, suppose that $C_0, \gamma_0 > 0$ are constants so that, for all $\gamma \geq \gamma_0$, $\bar{m}(\gamma) \leq C_0 \gamma \log \gamma$. (Such constants are guaranteed to exist by Theorem 3.) Further, suppose the*

number of scenarios $n \triangleq n(\gamma)$ is chosen as a function of γ , and that there exist constants $C_1, \gamma_1 > 0$, so that for all $\gamma \geq \gamma_1$, $n(\gamma) \geq C_1\gamma$. That is, n asymptotically grows at least linearly in γ . Then, for any $\epsilon, \delta > 0$, there exists $\gamma_2 > 0$ so that for all $\gamma \geq \gamma_2$,

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n m_i \geq (C_0 + \epsilon)\gamma \log \gamma \right) < \delta.$$

5.3. Optimal Non-Uniform Threshold Estimator

Theorems 2 and 3 and Corollary 2 allow a comparison between the UNIFORM estimator and the non-uniform THRESHOLD estimator. In particular, suppose $\hat{\alpha}_{m,n}$ is the UNIFORM estimate with n scenarios and m inner stage samples. As discussed in Section 3, when $m, n \rightarrow \infty$, this has asymptotic bias and variance

$$(17) \quad \mathbb{E} [\hat{\alpha}_{m,n} - \alpha] = \frac{\theta_c}{m} + O(m^{-3/2}), \quad \text{Var}(\hat{\alpha}_{m,n}) = \frac{\alpha(1-\alpha)}{n} + O(m^{-1}n^{-1}).$$

On the other hand, suppose that $\tilde{\alpha}_{\gamma,n}$ is the non-uniform THRESHOLD estimator with n scenarios and a threshold of γ . By Theorem 3, this estimator will employ, on average, $\bar{m} \triangleq \bar{m}(\gamma) = O(\gamma^{1+\epsilon})$ inner stage samples per scenario, for any positive ϵ . We can express the asymptotic bias and variance results of Theorem 2 and Corollary 2 as a function of n and \bar{m} by

$$(18) \quad \mathbb{E} [\tilde{\alpha}_{\gamma,n} - \alpha] = O(\bar{m}^{-2+\epsilon}), \quad \text{Var}(\tilde{\alpha}_{\gamma,n}) = \frac{\alpha(1-\alpha)}{n} + O(\bar{m}^{-2+\epsilon}n^{-1}),$$

for all positive ϵ .

Comparing (17) and (18), we see that, up to the dominant term, the two algorithms achieve the same asymptotic variance of order n^{-1} . This is consistent with the discussion in Section 3, which suggests that the asymptotic variance is determined by the randomness in scenario generation. This is exactly the same in the two algorithms. The inner stage sampling is different, however, and this results in a difference in bias for the estimators. Specifically, as a function of the average number of inner stage samples per scenario, the bias of the non-uniform THRESHOLD estimator decays approximately as the square of the bias of the UNIFORM estimator.

Given a total work budget of k (i.e., $mn \leq k$), we saw in Section 3 that the optimal UNIFORM estimator (in the sense of minimum MSE) would utilize a number of scenarios n of order $k^{1/3}$, a number of inner stage samples per scenario m of order $k^{2/3}$, and result in an MSE of order $k^{-2/3}$. For the non-uniform THRESHOLD estimator, from the results of Section 5.2, we can bound the MSE by

$$\mathbb{E} [(\tilde{\alpha}_{\gamma,n} - \alpha)^2] \leq \frac{\alpha(1-\alpha)}{n} + \frac{C}{\gamma^4},$$

for sufficiently large n and γ and an appropriate choice of the constant C . We can find a non-uniform THRESHOLD estimator with low MSE by minimizing this upper bound over choices of

(γ, n) , subject to an expected total work constraint. That is, we consider optimization problem

$$(19) \quad \begin{aligned} & \underset{\gamma, n}{\text{minimize}} && \frac{\alpha(1-\alpha)}{n} + \frac{C}{\gamma^4} \\ & \text{subject to} && \bar{m}(\gamma)n \leq k, \\ & && \gamma, n \geq 0. \end{aligned}$$

For any positive ϵ and given a work budget k , suppose we choose $\gamma^* \propto k^{1/5}$ and $n^* \propto k^{4/5-\epsilon}$. Then, we have that $\bar{m}(\gamma^*)n^* = O(k^{1-\epsilon} \log k) = o(k)$. Thus, for k sufficiently large, the expected total work will be less than k . Indeed, since (γ^*, n^*) satisfy the conditions of Theorem 4, for k sufficiently large the realized total work will also be less than k with high probability. This choice will result in MSE of $O(k^{-4/5+\epsilon})$. Hence, the optimal non-uniform THRESHOLD estimator converges at a *faster* rate than *any* uniform estimator. This is accomplished by generating more outer scenarios ($k^{4/5-\epsilon}$ vs. $k^{2/3}$) and perform less inner stage sampling on average in each scenario ($k^{1/5}$ vs. $k^{1/3}$) than is optimal in the uniform case.

6. Adaptive Allocation Algorithm

The non-uniform SEQUENTIAL estimator provides a way to determine the placement of inner stage samples across scenarios. The decision of how to allocate computational effort between generating more scenarios (i.e., the choice of n) and generating more inner samples across scenarios (i.e., the choice of \bar{m}) is unaddressed, however. The discussion in Section 5.3 suggests that, given a total work budget of k , one should asymptotically approximately choose $n \propto k^{4/5}$ and $\bar{m} \propto k^{1/5}$. However, the constants in these asymptotic expressions are unspecified. The choice of these constants may have an enormous impact on the practical performance of these algorithms. Note that the UNIFORM estimator faces the same problem — indeed, the optimal allocation (6) suggested by the analysis of Section 3 requires knowledge of the constant θ_c . It is not clear, in general, how to determine this constant.

In this section we will consider an *adaptive allocation* approach. This algorithm is a heuristic that estimates the optimal choice of \bar{m} and n at each point in time. It refines these estimates over the course of the simulation. The main idea of this approach is that, based on the results of Section 5, the variance is determined by the number of scenarios (n) and the bias squared is determined by the amount of inner sampling (\bar{m}). The adaptive algorithm estimates these quantities and then either increases the number of scenarios or increases the number of inner samples depending on whether the MSE is dominated by the variance or the biased squared.

Specifically, the ADAPTIVE estimator of Algorithm 4 proceeds as follows:

1. The simulation is initialized (lines 2–7) by generating n^0 scenarios with m^0 inner samples for each scenario.
2. The work budget of the simulation k is divided into $K \triangleq k/\tau_e$ intervals (or, epochs) of length

τ_e (note that we assume for simplicity of exposition that K is integral, and that the first epoch is only of length $\tau_e - n^0 m^0$ due to the initialization).

3. At the beginning of the ℓ th epoch (line 9), estimates are made for the bias squared and variance of the loss probability estimate, given the scenarios and samples that have been generated thus far. Specifically, given the loss probability estimate

$$\hat{\alpha} = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{\hat{L}_i \geq c\}},$$

the bias is approximated according to

$$(20) \quad \mathbb{E}[\hat{\alpha} - \alpha] \approx \hat{B} \triangleq \hat{\alpha} - \bar{\alpha},$$

where

$$\bar{\alpha} \triangleq \frac{1}{n} \sum_{i=1}^n \Phi \left(\frac{m_i(\hat{L}_i - c)}{\sigma_i} \right).$$

This approximation is based on a central limit theorem heuristic: in each scenario ω_i , when the number of samples m_i is large, each loss estimate \hat{L}_i can be approximated by a normal distribution with mean equal to $L(\omega_i)$ and with variance σ_i^2/m_i . Hence, given a fixed set of scenarios $\omega_1, \dots, \omega_n$, one might estimate the bias via

$$\mathbb{E}[\hat{\alpha} - \alpha] = \frac{1}{n} \sum_{i=1}^n \left\{ \mathbb{P}(\hat{L}_i \geq c) - \mathbb{I}_{\{L(\omega_i) \geq c\}} \right\} \approx \frac{1}{n} \sum_{i=1}^n \left\{ \Phi \left(\frac{m_i(L(\omega_i) - c)}{\sigma_i} \right) - \mathbb{I}_{\{L(\omega_i) \geq c\}} \right\}.$$

Since each true loss $L(\omega_i)$ is unknown in practice, we can approximate this with its realized estimate \hat{L}_i . This results in (20). By making a similar heuristic approximation for the variance, we arrive at the expression

$$(21) \quad \text{Var}(\hat{\alpha}) \approx \hat{V} \triangleq \frac{\bar{\alpha}(1 - \bar{\alpha})}{n}.$$

Note that the estimators (20) and (21) are meant only as heuristics. Better estimators may be possible and bias, in particular, is notoriously difficult to estimate. For our purposes, however, they only need to be accurate within orders of magnitude so as to allocate computational effort between inner samples and outer scenarios. We will see in the numerical results of Section 7 that, empirically, they suffice for this purpose.

4. Suppose there are n outer scenarios and an average of $\bar{m} \triangleq \frac{1}{n} \sum_{i=1}^n m_i$ inner samples per scenario at the beginning of the ℓ th epoch. From the results in Section 5, we expect the bias squared to decrease according to $\bar{m}^{-4+\epsilon}$ and the variance to decrease in proportion to n^{-1} . Then, assume that the number of scenarios and samples at the *end* of the ℓ th epoch is given by n' and \bar{m}' . We can estimate the bias squared at the end of the ℓ th epoch, as a function

of the bias estimate \hat{B} at the beginning, by $\hat{B}^2(\bar{m}/\bar{m}')^4$. Similarly, the variance at the end of the ℓ th epoch can be estimated by $\hat{V}(\frac{n}{n'})$.

Thus, at the beginning of the ℓ th epoch, we consider the following optimization problem:

$$(22) \quad \begin{aligned} & \underset{\bar{m}', n'}{\text{minimize}} && \hat{B}^2 \left(\frac{\bar{m}}{\bar{m}'} \right)^4 + \hat{V} \left(\frac{n}{n'} \right). \\ & \text{subject to} && \bar{m}' n' = \bar{m} n + \tau_e, \\ & && n \leq n' \leq n + \tau_e, \\ & && \bar{m}' \geq 0. \end{aligned}$$

This problem seeks to make a choice of (\bar{m}', n') that results in a minimal mean squared error at the end of the ℓ th epoch. The first constraint ensures that the total number of inner samples in the ℓ th epoch will equal the epoch length τ_e . The second constraint ensures that the number of scenarios at the end of the ℓ th epoch is at least the number of scenarios at the beginning, and increases by at most the length of the epoch.

The solution to (22) is given by

$$(23) \quad n' = \min \left\{ \max \left\{ \left(\frac{\hat{V} n}{4 \hat{B}^2 \bar{m}^4} (\bar{m} n + \tau_e)^4 \right)^{1/5}, n \right\}, n + \tau_e \right\}, \quad \bar{m}' = \frac{\bar{m} n + \tau_e}{n'}.$$

After obtaining the target number of scenarios n' (line 10), $n' - n$ additional scenarios are generated.

5. Over the course of the ℓ th epoch (lines 13–21) τ_e inner samples are generated. These are distributed to ensure that every scenario has at least m^0 inner samples in total (not per epoch). Once that is the case, inner samples are allocated myopically according to minimum error margin as in the SEQUENTIAL estimator.

7. Numerical Results

In this section we present numerical results that illustrate the benefits of non-uniform nested estimation. We begin in Section 7.1 by describing two settings for our numerical experiments. In Section 7.2, we compare the bias of the UNIFORM estimator and the non-uniform THRESHOLD and SEQUENTIAL estimators. In Section 7.3, we compare the MSE of a number of both implementable and idealized uniform and non-uniform estimators. Finally, in Section 7.4, we consider issues arising from the estimation of the variance of inner stage samples.

7.1. Experimental Setting

Our numerical experiments are set in the context of the following two examples: a portfolio with Gaussian cashflows, where both the outer stage scenarios and inner stage samples are generated

```

1: procedure ADAPTIVE( $m^0, n^0, \tau_e, k$ )
2:   generate scenarios  $\omega_1, \omega_2, \dots, \omega_{n^0}$ 
3:    $n \leftarrow n^0$ 
4:   for  $i \leftarrow 1$  to  $n^0$  do
5:     conditioned on scenario  $\omega_i$ , generate i.i.d. samples  $\hat{Z}_{i,1}, \dots, \hat{Z}_{i,m^0}$  of portfolio losses
6:      $m_i \leftarrow m^0$ 
7:   end for
8:   for  $\ell \leftarrow 1$  to  $\lceil k/\tau_e \rceil$  do
9:     estimate the current bias and variance by  $\hat{B}$  and  $\hat{V}$  from (20)–(21)
10:    determine a target number of scenarios by
11:
12:     
$$n' \leftarrow \left[ \min \left\{ \max \left\{ \left( \frac{\hat{V}n}{4\hat{B}^2\bar{m}^4} (\bar{m}n + \tau_e)^4 \right)^{1/5}, n \right\}, n + \tau_e \right\} \right]$$

13:    generate scenarios  $\omega_{n+1}, \dots, \omega_{n'}$ , set  $m_i \leftarrow 0$  for  $i = n + 1, \dots, n'$ 
14:     $n \leftarrow n'$ 
15:    while  $\sum_{i=1}^n m_i < \ell\tau_e$  do
16:      if  $\min_i m_i < m^0$  then
17:        set  $i^* \in \operatorname{argmin}_i m_i$ 
18:      else
19:        set  $i^* \in \operatorname{argmin}_i m_i |\hat{L}_i - c|/\sigma_i$ 
20:      end if
21:      generate one additional portfolio loss sample  $\hat{Z}_{i^*, m_{i^*}+1}$  in scenario  $\omega_{i^*}$ 
22:       $m_{i^*} \leftarrow m_{i^*} + 1$ 
23:    end while
24:  end for
25:  compute an estimate of the probability of a large loss,  $\hat{\alpha} \leftarrow \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\hat{L}_i \geq c\}}$ 
26:  return  $\hat{\alpha}$ 
27: end procedure

```

Algorithm 4: Estimate the probability of a large loss using an adaptive non-uniform nested simulation. This estimator employs a sequential algorithm to determine the placement of inner stage samples across scenarios, and adaptively decides the number of scenarios and inner samples to add by estimating the bias and variance. The parameters n^0 and m^0 are the initial number of scenarios and inner samples per scenario, respectively. The parameter τ_e is the epoch length. The parameter k is the total number of inner samples. Note that each standard deviation σ_i can be estimated in an online fashion over the course of the simulation, as is discussed in Section 7.4.

from normal distributions, and a put option example, where the portfolio consists of a single put option on an underlying asset whose price follows a geometric Brownian motion process. For both examples, we are interested in computing the probability of a loss. We consider loss thresholds corresponding to 10%, 1%, and 0.1% loss probabilities.

In the Gaussian example, we consider a portfolio with normally distributed risk factors and cashflows. This is the simplest setting in which to test any nested simulation procedure. Specifically, we consider a portfolio with value $X_0 = 0$ at time $t = 0$ and value $X_\tau(\omega) = \omega$ at the risk horizon τ . We assume that the real-valued risk factor $\omega \in \mathbb{R}$ is normally distributed with mean zero and standard deviation $\sigma_{\text{outer}}^2 = 1$. Then, the loss $L(\omega) = X_0 - X_\tau(\omega) = -\omega$ is a standard normal random variable. Given a scenario ω_i , each inner loss sample takes the form $\hat{Z}_{i,j} = -\omega_i + \sigma_{\text{inner}} W_{i,j}$, where $W_{i,j}$ is a standard normal random variable and $\sigma_{\text{inner}} = 5$ is the standard deviation of the inner stage samples.

In this case, given a loss threshold c , the probability of a loss exceeding c is given by $\alpha = \Phi(-c)$. We choose the values 1.282, 2.326, and 3.090 for the loss threshold c , corresponding to loss probabilities α of 10%, 1%, and 0.1%, respectively.

In the put option example, we assume that the portfolio consists of a long position in a single put option. This example is more complex since the portfolio cashflows are non-linear and follow highly skewed distributions which vary substantially across outer stage scenarios. Here, the underlying asset follows a geometric Brownian motion with an initial price of $S_0 = 100$. The drift of this process under the real-world distribution used in the outer stage of simulation is $\mu = 8\%$. The annualized volatility is $\sigma = 20\%$. The risk-free rate is $r = 3\%$. The strike of the put option is $K = 95$ and the maturity is $T = 0.25$ years (i.e., three months). The risk horizon is $\tau = 1/52$ years (i.e., one week). With these parameters, the initial value of the put is $X_0 = 1.669$ given by the Black-Scholes formula.

Denote by $S_\tau(\omega)$ the underlying asset price at the risk horizon τ . This random variable is generated according to $S_\tau(\omega) \triangleq S_0 e^{(\mu - \sigma^2/2)\tau + \sigma\sqrt{\tau}\omega}$, where the real-valued risk-factor ω is a standard normal random variable. The portfolio loss at the risk horizon τ is given by

$$L(\omega) = X_0 - \mathbf{E} \left[e^{-r(T-\tau)} \max(K - S_T(\omega, W), 0) \mid \omega \right],$$

where the expectation is taken over the random variable W , which is an independently distributed standard normal, and $S_T(\omega, W)$ is given by

$$S_T(\omega, W) \triangleq S_\tau(\omega) e^{(r - \sigma^2/2)(T-\tau) + \sigma\sqrt{T-\tau}W}.$$

Note that, given a fixed value of ω and a standard normal W , the random variable $S_T(\omega, W)$ is distributed according to the risk-neutral distribution of underlying asset price at the option maturity T , conditional on asset price $S_\tau(\omega)$ at the risk horizon τ . Given an outer scenario ω_i , each

inner loss sample takes the form

$$\hat{Z}_{i,j} = X_0 - e^{-r(T-\tau)} \max(K - S_T(\omega_i, W_{i,j}), 0),$$

where $W_{i,j}$ is an independent standard normal random variable. Notice that outer stage scenarios are generated using the real-world distribution governed by the drift μ , while inner stage scenarios used to generate future put option prices are generated using the risk-neutral distribution governed by the drift r .

It is not difficult to see that the loss $L(\omega)$ is strictly increasing in the risk factor ω . Hence, the probability of a loss exceeding a threshold c can be computed according to $\alpha = \mathbb{P}(L \geq c) = \mathbb{P}(\omega \geq \omega^*) = \Phi(-\omega^*)$, where ω^* is the unique solution to $L(\omega^*) = c$. We choose the values 0.859, 1.221, and 1.390 for the loss threshold c , corresponding to loss probabilities α of 10%, 1%, and 0.1%, respectively.

7.2. Bias Comparison

As established in Section 5, the advantage of non-uniform inner stage sampling relative to uniform sampling is that, for the same total quantity of inner samples, a lower bias is attained. In this section, we numerically compare the UNIFORM estimator and the non-uniform THRESHOLD and SEQUENTIAL estimators on the basis of bias.

For this purpose, we generate a fixed sequence $\omega_1, \dots, \omega_n$ of $n = 10,000$ outer stage scenarios. In order to eliminate any noise in our comparison due to randomness in scenario generation, we choose the scenarios in a deterministic and stratified manner, so that $\mathbb{P}(\omega \leq \omega_i) = i/(n+1)$, for all $1 \leq i \leq n$. Given the stratified scenarios, we numerically compute the bias of each estimator, measured over 1,000 independent trials, as the total number of inner stage samples (or, the work budget) is varied from $k = 20,000$ to $k = 4,000,000$. In the case of the UNIFORM estimator, this is accomplished by varying the number of inner stage samples per scenario from $m = 2$ to $m = 400$. For the non-uniform SEQUENTIAL estimator, this is accomplished by using $m^0 = 2$ initial inner samples per scenario, and then varying the average number of inner stage samples per scenario from $\bar{m} = 2$ to $\bar{m} = 400$. In the case of the non-uniform THRESHOLD estimator, the threshold parameter γ was varied over the interval $(5 \times 10^{-5}, 2 \times 10^{-1})$ and the expected total number of inner stage samples was plotted (averaged over the independent trials). This range of γ was experimentally chosen so that the range of expected total inner samples for the THRESHOLD algorithm coincided with the range of total inner samples for the other algorithms.

The results for both the Gaussian example and the put option example with $\alpha = 1\%$ are plotted in Figure 5. In both cases, the non-uniform THRESHOLD and SEQUENTIAL estimators exhibit a lower bias than the UNIFORM estimator, given the same work budget. Further, for the UNIFORM estimator, the results are consistent with the bias decreasing with order k^{-1} , as suggested by Theorem 1. For the non-uniform THRESHOLD estimator, the results are consistent⁸ with the

⁸The exact rates of decay (i.e., the asymptotic slopes in Figure 5) are challenging to accurately estimate numerically.

bias decreasing according to $k^{-2+\epsilon}$ for any positive ϵ , as suggested by the theory presented in Section 5. Note that the performance of the THRESHOLD and SEQUENTIAL estimators is largely indistinguishable. This strongly suggests that our theoretical analysis of the rate of convergence of THRESHOLD estimator in Section 5 provides a good proxy for the rate of convergence of the SEQUENTIAL estimator.


Figure 6 gives some qualitative insight into the inner sampling behavior of the non-uniform SEQUENTIAL estimator. Here, we have plotted the number of inner samples (averaged across the 1,000 independent trials) at a scenario against the loss in the scenario. Here, the amount of inner sampling employed by the SEQUENTIAL varies over two orders of magnitude across scenarios, with much more sampling is taking place close to the loss threshold c than far away from it.

7.3. MSE Comparison

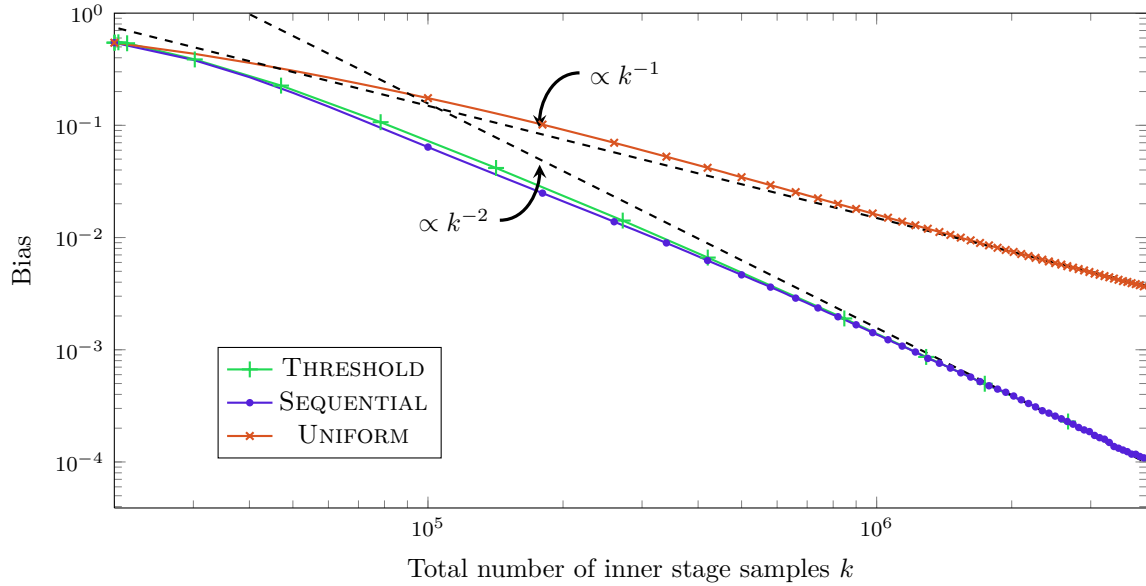
In this section, we will provide an overall comparison of the MSE achieved by various uniform and non-uniform estimators, given a fixed computational budget of k inner stage samples. We consider each of the following estimators:

- **Optimal uniform.** This is the UNIFORM estimator with parameters chosen optimally, as in Section 5. The simulation budget is allocated according to $m = k^{1/3}/\beta^*$ and $n = \beta^*k^{2/3}$, where the constant β^* , given in (6), is chosen to minimize MSE. Note that, in general, it is not clear how to determine the value β^* given the problem parameters. For both the Gaussian and put option examples here, we are able to use closed form expressions for the probability distribution of losses to exactly compute this constant.
- $1/3 : 2/3$ **uniform.** This is the UNIFORM estimator with $m = k^{1/3}$ and $n = k^{2/3}$. Based on the analysis in Section 3, this estimator has MSE that decays with same order ($k^{-2/3}$) as the optimal uniform estimator, but with a suboptimal constant. This is meant to illustrate the case where the constant β^* of the optimal uniform estimator is unknown, and an arbitrary choice of constant ($\beta^* = 1$) is made.
- **Optimal sequential.** This is the SEQUENTIAL non-uniform sampling estimator where n is chosen optimally to minimize MSE. Here, $m^0 = 2$ initial samples were used. The parameters (\bar{m}, n) controlling the average number of inner samples and the number of scenarios were varied over choices satisfying the simulation work budget, i.e., $\bar{m}n = k$. The choice which resulted in minimum MSE was selected. The optimal sequential estimator is an idealized algorithm meant to capture the best possible performance than can be achieved using the SEQUENTIAL estimator.
- **Adaptive.** This is the ADAPTIVE estimator of Section 6, which utilizes sequential non-uniform sampling and adaptively allocates computational effort between outer stage scenarios

This is because it is not computationally feasible to compute the estimators over many orders of magnitude of k . The results in Figure 5 are not intended as numerical “proof” of a particular rate of convergence, but are rather intended to illustrate that the numerical convergence is consistent with our earlier theoretical analyses.




(a) Gaussian, $\alpha = 1\%$




(b) Put option, $\alpha = 1\%$

Figure 5: The vertical axis shows the bias in absolute terms, i.e., the absolute value of difference between the estimated loss probability and the true loss probability α , as a function of the total number of inner stage samples. In the case of the THRESHOLD algorithm, the expected total number of samples is shown. A set of $n = 10,000$ stratified outer scenarios was used. The bias of the non-uniform THRESHOLD and SEQUENTIAL estimators is consistent with the predicted theoretical decay of $k^{-2+\epsilon}$, for any positive ϵ . Similarly, the bias of the UNIFORM estimator is consistent with the predicted theoretical decay of k^{-1} .



(a) Gaussian, $\alpha = 1\%$



(b) Put option, $\alpha = 1\%$

Figure 6: The number of inner stage samples as a function of the loss in each scenario, averaged over 1,000 trials. Here, $k = 4,000,000$ inner stage samples are distributed across $n = 10,000$ stratified scenarios. The UNIFORM estimator employs $m = 400$ inner samples for each scenario. The non-uniform SEQUENTIAL estimator varies the number of samples over two orders of magnitude, and employs many more samples close to the loss threshold c .

and inner stage samples. Here, $n^0 = 500$ initial scenarios were used, with $m^0 = 2$ initial inner samples per scenario. An epoch length of $\tau_e = 100,000$ was used.

- **Adaptive ($\hat{\sigma}_i$).** This is a variation of the ADAPTIVE estimator in which the variance of inner samples is estimated. This will be discussed shortly in Section 7.4.

The numerical results for the six test cases (the Gaussian and put option examples, each with thresholds corresponding to three different loss probabilities) using the five estimators are summarized in Table 1. In all cases, a computational budget of $k = 4,000,000$ inner stage samples was used. The results in each case are computed over 1,000 independent trials.

The numerical results in Table 1 can be interpreted naturally through a series of pairwise comparisons, as follows:

- **Optimal uniform vs. $1/3:2/3$ uniform.** These are both asymptotically optimal UNIFORM estimators; they differ only by the choice of constant β^* . The practical performance of these two estimators, however, is dramatically different. This highlights the sensitivity of the UNIFORM estimator in practice to the choice of constant. Note that computing the constant β^* , as given in (6), requires knowledge of the constant θ_c , defined in (5). It is not clear how to estimate this constant in practice, and this constant may vary dramatically across different problem instances.
- **Optimal uniform vs. optimal sequential.** These represent the best possible performance that can be achieved by the UNIFORM and SEQUENTIAL estimators. Neither of these estimators is implementable in practice — the former because it depends on a parameter that cannot be readily determined from the problem data, the latter because it requires exploration over the choice of parameters. However, by contrasting them we can see a comparison of uniform and non-uniform sampling on an equal footing. This comparison clearly illustrates benefits of non-uniform sampling. In every test case, the optimal sequential estimator has the lowest MSE. The MSE improvement relative to the optimal uniform estimator is between a factor of 4 and 10. This improvement is greatest when estimating loss probabilities that are rare (e.g., the $\alpha = 0.1\%$ case).

Further, note that the optimal sequential estimator employs many fewer inner stage samples and many more outer stage scenarios. This is consistent with the theory developed in Section 5 and the experiments in Section 7.2. The optimal sequential estimator is able to achieve a low bias with fewer inner stage samples, hence it can employ more scenarios with the same computational budget.

- **Optimal sequential vs. adaptive sequential.** The optimal sequential estimator relies on a brute force optimization over the parameters choosing the number of inner samples and outer scenarios; this is not feasible in practice. On the other hand, the adaptive sequential makes this choice dynamically over the course of the simulation and thus is implementable

		n	\bar{m}	variance	bias ²	MSE	MSE std. err.	MSE norm.
Gaussian $\alpha = 10\%$	$1/3:2/3$ uniform	25,199	159	$4.0 \cdot 10^{-6}$	$2.8 \cdot 10^{-4}$	$2.9 \cdot 10^{-4}$	$2.1 \cdot 10^{-6}$	35.4
	optimal uniform	4,499	889	$2.1 \cdot 10^{-5}$	$8.6 \cdot 10^{-6}$	$3.0 \cdot 10^{-5}$	$1.2 \cdot 10^{-6}$	3.7
	adaptive ($\hat{\sigma}_i$)	14,968	281	$7.0 \cdot 10^{-6}$	$2.7 \cdot 10^{-6}$	$9.7 \cdot 10^{-6}$	$4.7 \cdot 10^{-7}$	1.2
	adaptive	12,802	321	$7.2 \cdot 10^{-6}$	$1.5 \cdot 10^{-6}$	$8.6 \cdot 10^{-6}$	$3.9 \cdot 10^{-7}$	1.0
	optimal sequential	12,395	323	$6.8 \cdot 10^{-6}$	$1.4 \cdot 10^{-6}$	$8.2 \cdot 10^{-6}$	$3.7 \cdot 10^{-7}$	1.0
Gaussian $\alpha = 1\%$	$1/3:2/3$ uniform	25,199	159	$6.1 \cdot 10^{-7}$	$2.8 \cdot 10^{-5}$	$2.8 \cdot 10^{-5}$	$2.6 \cdot 10^{-7}$	60.9
	optimal uniform	5,089	786	$2.3 \cdot 10^{-6}$	$1.0 \cdot 10^{-6}$	$3.3 \cdot 10^{-6}$	$1.5 \cdot 10^{-7}$	7.2
	adaptive ($\hat{\sigma}_i$)	16,177	250	$7.0 \cdot 10^{-7}$	$3.7 \cdot 10^{-9}$	$7.0 \cdot 10^{-7}$	$3.1 \cdot 10^{-8}$	1.5
	adaptive	16,118	251	$7.1 \cdot 10^{-7}$	$4.1 \cdot 10^{-9}$	$7.2 \cdot 10^{-7}$	$3.1 \cdot 10^{-8}$	1.6
	optimal sequential	30,860	130	$3.5 \cdot 10^{-7}$	$1.1 \cdot 10^{-7}$	$4.6 \cdot 10^{-7}$	$1.8 \cdot 10^{-8}$	1.0
Gaussian $\alpha = 0.1\%$	$1/3:2/3$ uniform	25,199	159	$8.2 \cdot 10^{-8}$	$1.1 \cdot 10^{-6}$	$1.2 \cdot 10^{-6}$	$1.9 \cdot 10^{-8}$	48.0
	optimal uniform	7,788	514	$1.7 \cdot 10^{-7}$	$7.9 \cdot 10^{-8}$	$2.5 \cdot 10^{-7}$	$1.3 \cdot 10^{-8}$	10.0
	adaptive ($\hat{\sigma}_i$)	30,798	132	$3.5 \cdot 10^{-8}$	$4.7 \cdot 10^{-10}$	$3.5 \cdot 10^{-8}$	$1.6 \cdot 10^{-9}$	1.4
	adaptive	30,628	132	$3.8 \cdot 10^{-8}$	$5.0 \cdot 10^{-10}$	$3.8 \cdot 10^{-8}$	$3.2 \cdot 10^{-9}$	1.5
	optimal sequential	56,686	71	$1.8 \cdot 10^{-8}$	$6.5 \cdot 10^{-9}$	$2.5 \cdot 10^{-8}$	$1.1 \cdot 10^{-9}$	1.0
Put option $\alpha = 10\%$	$1/3:2/3$ uniform	25,199	159	$4.1 \cdot 10^{-6}$	$5.1 \cdot 10^{-4}$	$5.1 \cdot 10^{-4}$	$2.9 \cdot 10^{-6}$	58.6
	optimal uniform	5,095	785	$1.9 \cdot 10^{-5}$	$2.4 \cdot 10^{-5}$	$4.2 \cdot 10^{-5}$	$1.6 \cdot 10^{-6}$	4.8
	adaptive ($\hat{\sigma}_i$)	6,671	601	$1.5 \cdot 10^{-5}$	$4.8 \cdot 10^{-6}$	$2.0 \cdot 10^{-5}$	$9.2 \cdot 10^{-7}$	2.3
	adaptive	7,325	547	$1.3 \cdot 10^{-5}$	$2.1 \cdot 10^{-7}$	$1.4 \cdot 10^{-5}$	$6.2 \cdot 10^{-7}$	1.6
	optimal sequential	12,395	323	$7.3 \cdot 10^{-6}$	$1.5 \cdot 10^{-6}$	$8.7 \cdot 10^{-6}$	$3.8 \cdot 10^{-7}$	1.0
Put option $\alpha = 1\%$	$1/3:2/3$ uniform	25,199	159	$7.8 \cdot 10^{-7}$	$9.4 \cdot 10^{-5}$	$9.5 \cdot 10^{-5}$	$5.4 \cdot 10^{-7}$	141.8
	optimal uniform	3,143	1,273	$3.8 \cdot 10^{-6}$	$1.2 \cdot 10^{-6}$	$5.0 \cdot 10^{-6}$	$2.1 \cdot 10^{-7}$	7.5
	adaptive ($\hat{\sigma}_i$)	10,085	401	$1.2 \cdot 10^{-6}$	$2.0 \cdot 10^{-7}$	$1.4 \cdot 10^{-6}$	$6.2 \cdot 10^{-8}$	2.1
	adaptive	9,992	405	$1.1 \cdot 10^{-6}$	$1.7 \cdot 10^{-8}$	$1.1 \cdot 10^{-6}$	$4.8 \cdot 10^{-8}$	1.6
	optimal sequential	19,558	205	$5.4 \cdot 10^{-7}$	$1.5 \cdot 10^{-7}$	$6.9 \cdot 10^{-7}$	$3.0 \cdot 10^{-8}$	1.0
Put option $\alpha = 0.1\%$	$1/3:2/3$ uniform	25,199	159	$1.5 \cdot 10^{-7}$	$8.1 \cdot 10^{-6}$	$8.2 \cdot 10^{-6}$	$7.2 \cdot 10^{-8}$	174.5
	optimal uniform	2,570	1,556	$4.4 \cdot 10^{-7}$	$3.9 \cdot 10^{-8}$	$4.8 \cdot 10^{-7}$	$2.7 \cdot 10^{-8}$	10.2
	adaptive ($\hat{\sigma}_i$)	14,884	274	$1.1 \cdot 10^{-7}$	$1.8 \cdot 10^{-8}$	$1.3 \cdot 10^{-7}$	$9.0 \cdot 10^{-9}$	2.8
	adaptive	14,384	284	$9.2 \cdot 10^{-8}$	$5.6 \cdot 10^{-10}$	$9.2 \cdot 10^{-8}$	$1.4 \cdot 10^{-8}$	2.0
	optimal sequential	26,508	151	$3.9 \cdot 10^{-8}$	$8.0 \cdot 10^{-9}$	$4.7 \cdot 10^{-8}$	$2.3 \cdot 10^{-9}$	1.0

Table 1: Numerical results for five estimation algorithms over six test cases (the Gaussian and put option examples, each with thresholds corresponding to three different loss probabilities). The results are computed over 1,000 independent trials, each with a total simulation budget of $k = 4,000,000$. The results reported include the number of outer stage scenarios (n) and the average number of inner stage samples per scenario (\bar{m}) employed by each estimator, as well as the variance, the bias squared, the mean squared error (MSE), and the standard error of the MSE for each estimator. The last column contains MSE results normalized relative to the optimal sequential estimator.

in practice. Comparing these two methods illustrates how much of the benefit of the optimal sequential method can be achieved in practice.

Across our experiments, adaptive sequential estimator achieves an MSE between 1 and 2 times that of the optimal sequential estimator. In some cases, the adaptive estimator overestimates the true bias and uses too many inner stage samples compared to the optimal allocation. This suggests that there is modest room for improvement in the ADAPTIVE procedure for allocating computational effort between inner and outer stages.

7.4. Variance Estimation

The ADAPTIVE algorithm requires the value of σ_i , the standard deviation of the inner stage loss samples $\hat{Z}_{i,1}, \hat{Z}_{i,2}, \dots$ in scenario ω_i . In practice, σ_i will not be known. However, one can imagine many variations of the ADAPTIVE algorithm where each σ_i is estimated over the course of the estimation algorithm.

One such variation replaces each σ_i in the ADAPTIVE algorithm with the estimate

$$(24) \quad \hat{\sigma}_i \triangleq \frac{m_i}{m_i + b} \tilde{\sigma}_i + \frac{b}{m_i + b} \bar{\sigma}.$$

Here, we define

$$\tilde{\sigma}_i \triangleq \left[\frac{1}{m_i - 1} \sum_{j=1}^{m_i} (\hat{Z}_{i,j} - \hat{L}_i)^2 \right]^{1/2}$$

to be the sample standard deviation of the inner stage loss samples generated in scenario ω_i , and $\bar{\sigma} \triangleq \frac{1}{n} \sum_{i=1}^n \tilde{\sigma}_i$ to be the overall average of all such sample standard deviations. This procedure balances an ensemble estimate with a local estimate so that the estimated standard deviations can be generated more reliably, especially when there are a small number of inner stage samples at a given scenario. For $b = 0$, the procedure corresponds to the usual sample standard deviation estimator. For large values of b , the ensemble estimate is given a larger weight.

Numerical results for an adaptive estimator using this procedure for estimating σ_i , with $b = 5$, are given in Table 1. To avoid a prohibitive computational burden, we only update the average $\bar{\sigma}$ at the end of each specific epoch.⁹ The results show that there is a modest to no loss in performance when the estimated $\hat{\sigma}_i$ is used in place of the true σ_i .

8. Conclusion

Two-level nested simulation can provide a more realistic assessment of financial risk, but with a considerable computational cost. In this paper we propose a nested sequential simulation procedure

⁹Numerically stable and efficient algorithms are available for updating sample variance calculations (see, e.g., Chan et al., 1983). These would allow for rapid calculation of each $\tilde{\sigma}_i$ in an online fashion. However, once $\bar{\sigma}$ is updated, *every* $\hat{\sigma}_i$ will change. This will necessitate rebuilding the priority queue data structure for the scenarios, and may require order n time.

which significantly reduces the computational burden. The savings are achieved by using a non-uniform inner sampling procedure which allocates more resources where the effect on the risk estimation is the greatest, which in turn allows relatively more effort to be devoted to the generation of outer scenarios. The combined effect produces a risk estimator which converges at a faster rate to the true value. In numerical experiments, mean squared error was reduced by factors ranging from four to over one hundred.

The sequential estimation procedure can be combined with previous research on variance reduction for the outer stage scenario generation to achieve further computational savings. The algorithms and results were presented in the context of estimating the probability of a large loss, but it may be possible to apply similar ideas to develop non-uniform algorithms for other risk measures. This remains an open area for future research.

References

- P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9: 203–228, 2000.
- T. F. Chan, G. H. Golub, and R. J. LeVeque. Algorithms for computing the sample variance: Analysis and recommendations. *The American Statistician*, 37(3):242–247, 1983.
- C.-H. Chen and L. Lee. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. World Scientific, 2010.
- T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 2nd edition, 2002.
- M. Crouhy, D. Galai, and R. Mark. *Risk Management*. McGraw-Hill, New York, 2000.
- H. Föllmer and A. Schied. Robust preferences and convex measures of risk. In *Advances in Finance and Stochastics*, pages 39–56, Berlin, 2002. Springer-Verlag.
- P. I. Frazier, W. B. Powell, and S. Dayanik. A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, 47(5):2410–2439, 2008.
- P. Glasserman, P. Heidelberger, and P. Shahabuddin. Variance reduction techniques for estimating value-at-risk. *Management Science*, 46:1349–1364, 2000.
- P. Glasserman, P. Heidelberger, and P. Shahabuddin. Portfolio value-at-risk with heavy-tailed risk factors. *Mathematical Finance*, 12:239–270, 2002.
- M. B. Gordy and S. Juneja. Nested simulation in portfolio risk measurement. FEDS 2008-21, Federal Reserve Board, April 2008.
- M. B. Gordy and S. Juneja. Nested simulation in portfolio risk management. *Management Science*, 56(10): 1833–1848, October 2010.
- A. Gut. On the moments and limit distributions of some first passage times. *The Annals of Probability*, 2 (2):277–308, April 1974.
- L. J. Hong and S. Juneja. Estimating the mean of a non-linear function of conditional expectation. In *Proceedings of the 2009 Winter Simulation Conference*, pages 1223–1236, Piscataway, NJ, 2009. IEEE Press.

- P. Jorion. *Value at Risk*. McGraw-Hill, New York, 2006.
- S.-H. Kim and B. L. Nelson. Selecting the best system. In *Elsevier Handbooks in Operations Research and Management Science: Simulation*, pages 501–534, Oxford, UK, 2005. Elsevier.
- H. Lan, B. L. Nelson, and J. Staum. Two-level simulations for risk management. In *INFORMS Simulation Society Research Workshop*, pages 102–107, 2007.
- H. Lan, B. L. Nelson, and J. Staum. A confidence interval procedure for expected shortfall risk measurement via two-level simulation. To appear in *Operations Research*, 2008.
- S.-H. Lee. *Monte Carlo Computation of Conditional Expectation Quantiles*. PhD thesis, Stanford University, 1998.
- S.-H. Lee and P. W. Glynn. 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, July 2003.
- V. Lesnevski, B. L. Nelson, and J. Staum. Simulation of coherent risk measures. In *Proceedings of the 2004 Winter Simulation Conference*, pages 1579–1585, Piscataway, NJ, 2004. IEEE Press.
- V. Lesnevski, B. L. Nelson, and J. Staum. Simulation of coherent risk measures based on generalized scenarios. *Management Science*, 53(11):1756–1769, November 2007.
- M. Liu and J. Staum. Estimating expected shortfall with stochastic kriging. Working paper, Northwestern University, 2009.
- M. Liu, B. L. Nelson, and J. Staum. An efficient simulation procedure for point estimation of expected shortfall. Working paper, Northwestern University, 2008.
- G. Lorden. On excess over the boundary. *Annals of Mathematical Statistics*, 41:520–527, 1970.
- A. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton University Press, Princeton, New Jersey, 2006.
- W. E. Pruitt. The growth of random walks and Levy processes. *Annals of Probability*, 9:948–956, 1981.
- R. T. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26:443–471, 2002.
- D. Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. Springer, New York, 1985.
- Y. Sun, D. W. Apley, and J. Staum. $1^{1/2}$ -level simulation for estimating the variance of a conditional expectation. Working paper, Northwestern University, 2009.

A. Proofs

In this section, we provide proofs for Theorems 2, 3, and 4 of Section 5.2, which analyze the performance of the THRESHOLD estimator.

A.1. Preliminaries

Consider the THRESHOLD estimator with n scenarios and a threshold parameter γ . Each scenario ω_i has inner loss samples $\hat{Z}_{i,1}, \hat{Z}_{i,2}, \dots$, which, by Assumption 2, are i.i.d. normal random variables with mean $L(\omega_i)$ and standard deviation $\sigma(\omega_i)$. The estimator will generate m_i inner stage samples in this scenario, with

$$(25) \quad m_i = \inf \left\{ m > 0 : \left| S_m^{(i)} \right| \geq \gamma \right\}.$$

Here, for $m \geq 0$, the partial sum $S_m^{(i)}$ is defined by $S_m^{(i)} \triangleq \sum_{j=1}^m (\hat{Z}_{i,j} - c)/\sigma(\omega_i)$. Each term in this partial sum has mean $\mu(\omega_i) \triangleq (L(\omega_i) - c)/\sigma(\omega_i)$. By considering these partial sums over all n scenarios, the THRESHOLD estimator can be written as $\tilde{\alpha}_{\gamma,n} = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{S_{m_i}^{(i)} \geq \gamma\}}$. We are interested in, as $\gamma \rightarrow \infty$, the asymptotic behavior of the bias,

$$(26) \quad \bar{b}(\gamma) \triangleq \mathbf{E}[\tilde{\alpha}_{\gamma,n} - \alpha] = \mathbf{E} \left[\mathbb{I}_{\{S_{m_i}^{(i)} \geq \gamma\}} - \mathbb{I}_{\{L(\omega_i) \geq c\}} \right],$$

and the expected number of inner stage samples per scenario,

$$(27) \quad \bar{m}(\gamma) \triangleq \mathbf{E}[m_i].$$

Now, given $\mu \in \mathbb{R}$, define \mathbf{P}_μ to be a probability measure so that, under \mathbf{P}_μ , the random variables Y_1, Y_2, \dots are a collection of i.i.d. normal random variables with mean μ and unit variance. For each $m \geq 0$, define the partial sum $S_m \triangleq \sum_{j=1}^m Y_j$. It follows from Assumption 2 that, if $\mu = \mu(\omega_i)$, then S_m has the same distribution as $S_m^{(i)}$. Define $m(\gamma) \triangleq \inf \{m > 0 : |S_m| \geq \gamma\}$. From (25) and (27), we have that $\bar{m}(\gamma)$, the expected number of inner stage samples for the THRESHOLD estimator, satisfies

$$(28) \quad \bar{m}(\gamma) = \int \mathbf{E}_\mu[m(\gamma)] p(\mu) d\mu.$$

Here, \mathbf{E}_μ denotes expectation under the distribution μ . Similarly, we can define

$$b_+(\gamma) \triangleq \mathbb{I}_{\{m(\gamma) < \infty \text{ and } S_{m(\gamma)} \geq \gamma\}}, \quad b_-(\gamma) \triangleq -\mathbb{I}_{\{m(\gamma) < \infty \text{ and } S_{m(\gamma)} \leq -\gamma\}},$$

$$b(\gamma) \triangleq b_-(\gamma) \mathbb{I}_{\{\mu \geq 0\}} + b_+(\gamma) \mathbb{I}_{\{\mu < 0\}}.$$

Then, from (26), we have that $\bar{b}(\gamma)$, the bias of the THRESHOLD estimator, satisfies

$$(29) \quad \bar{b}(\gamma) = \int \mathbf{E}_\mu[b(\gamma)] p(\mu) d\mu.$$

Finally, by Assumption 2, define $\delta \in (0, 1)$ so that p is continuously differentiable over the

interval $[-\delta, \delta]$, and set

$$(30) \quad U_0 \triangleq \max_{|\mu| \leq \delta} |p(\mu)|, \quad U_1 \triangleq \max_{|\mu| \leq \delta} |p'(\mu)|.$$

A.2. Asymptotic Bias

The asymptotic bias result of Theorem 2 is that, as $\gamma \rightarrow \infty$, $\bar{b}(\gamma) = O(\gamma^{-2})$. We will establish this via a careful analysis of (29). In particular, consider the decomposition

$$(31) \quad \begin{aligned} |\bar{b}(\gamma)| &\leq \left| \int_{|\mu| > \delta} \mathbf{E}_\mu[b(\gamma)]p(\mu) d\mu \right| + \left| \int_{|\mu| \leq \delta} \mathbf{E}_\mu[b(\gamma)]p(\mu) d\mu \right| \\ &\leq \left| \int_{|\mu| > \delta} \mathbf{E}_\mu[b(\gamma)]p(\mu) d\mu \right| + \left| \int_{|\mu| \leq \delta} \mathbf{E}_\mu[b(\gamma)]p(0) d\mu \right| + \left| \int_{|\mu| \leq \delta} \mathbf{E}_\mu[b(\gamma)]\mu p'(\varsigma(\mu)) d\mu \right|. \end{aligned}$$

Here, using Assumption 2, we have applied Taylor's theorem, and ς is a function with $|\varsigma(\mu)| \leq \delta$ for all $\mu \in [-\delta, \delta]$. By symmetry, for any μ , we have that $\mathbf{E}_\mu[b(\gamma)] = -\mathbf{E}_{-\mu}[b(\gamma)]$. Then,

$$(32) \quad \begin{aligned} |\bar{b}(\gamma)| &\leq \left| \int_{|\mu| > \delta} \mathbf{E}_\mu[b(\gamma)]p(\mu) d\mu \right| + \left| \int_{|\mu| \leq \delta} \mathbf{E}_\mu[b(\gamma)]\mu p'(\varsigma(\mu)) d\mu \right| \\ &\leq \left| \int_{|\mu| > \delta} \mathbf{E}_\mu[b(\gamma)]p(\mu) d\mu \right| + U_1 \int_{|\mu| \leq \delta} |\mathbf{E}_\mu[b(\gamma)]\mu| d\mu. \end{aligned}$$

Theorem 2 will follow by applying Lemmas 1 and 2, established below, to (32).

We begin with a preliminary proposition:

Proposition 1. *If $\mu < 0$, then*

$$e^{2\mu\gamma} \mathbf{E}_{-\mu} \left[(1 + 2\mu(S_{m(\gamma)} - \gamma)) \mathbb{I}_{\{S_{m(\gamma)} \geq \gamma\}} \right] \leq \mathbf{P}_\mu(S_{m(\gamma)} \geq \gamma) \leq e^{2\mu\gamma} \mathbf{P}_{-\mu}(S_{m(\gamma)} \geq \gamma).$$

If $\mu > 0$, then

$$e^{-2\mu\gamma} \mathbf{E}_{-\mu} \left[(1 + 2\mu(S_{m(\gamma)} + \gamma)) \mathbb{I}_{\{S_{m(\gamma)} \leq -\gamma\}} \right] \leq \mathbf{P}_\mu(S_{m(\gamma)} \leq -\gamma) \leq e^{-2\mu\gamma} \mathbf{P}_{-\mu}(S_{m(\gamma)} \leq -\gamma).$$

Proof. Consider the case where $\mu < 0$. Let F_μ denote the $N(\mu, 1)$ distribution. Note that the Radon-Nikodym derivative between the F_μ and $F_{-\mu}$ is given by $e^{2\mu y}$. Then,

$$(33) \quad \mathbf{P}_\mu(S_{m(\gamma)} \geq \gamma) = \mathbf{E}_\mu \left[\mathbb{I}_{\{S_{m(\gamma)} \geq \gamma\}} \right] = \mathbf{E}_{-\mu} \left[e^{2\mu S_{m(\gamma)}} \mathbb{I}_{\{S_{m(\gamma)} \geq \gamma\}} \right] = e^{2\mu\gamma} \mathbf{E}_{-\mu} \left[e^{2\mu(S_{m(\gamma)} - \gamma)} \mathbb{I}_{\{S_{m(\gamma)} \geq \gamma\}} \right].$$

For $x > 0$, we have that $1 - x \leq e^{-x} \leq 1$. Thus,

$$(1 + 2\mu(S_{m(\gamma)} - \gamma)) \mathbb{I}_{\{S_{m(\gamma)} \geq \gamma\}} \leq e^{2\mu(S_{m(\gamma)} - \gamma)} \mathbb{I}_{\{S_{m(\gamma)} \geq \gamma\}} \leq \mathbb{I}_{\{S_{m(\gamma)} \geq \gamma\}}.$$

The result follows after taking an expectation with respect to $\mathbf{P}_{-\mu}$, and applying (33). The case

where $\mu > 0$ is handled similarly. ■

Lemma 1. As $\gamma \rightarrow \infty$,

$$\int_{|\mu|>\delta} \mathbb{E}_\mu[b(\gamma)]p(\mu) d\mu = o(\gamma^{-2}).$$

Proof. Note that

$$\begin{aligned} \left| \int_{|\mu|>\delta} \mathbb{E}_\mu[b(\gamma)]p(\mu) d\mu \right| &= \left| \int_\delta^\infty \mathbb{E}_\mu[b_-(\gamma)]p(\mu) d\mu + \int_{-\infty}^{-\delta} \mathbb{E}_\mu[b_+(\gamma)]p(\mu) d\mu \right| \\ &\leq \int_\delta^\infty \mathbb{E}_\mu[|b_-(\gamma)|]p(\mu) d\mu + \int_{-\infty}^{-\delta} \mathbb{E}_\mu[|b_+(\gamma)|]p(\mu) d\mu \\ &\leq \int_\delta^\infty \mathbb{P}_\mu(S_{m(\gamma)} \leq -\gamma)p(\mu) d\mu + \int_{-\infty}^{-\delta} \mathbb{P}_\mu(S_{m(\gamma)} \geq \gamma)p(\mu) d\mu. \end{aligned}$$

By Proposition 1,

$$\begin{aligned} \left| \int_{|\mu|>\delta} \mathbb{E}_\mu[b(\gamma)]p(\mu) d\mu \right| &\leq \int_\delta^\infty e^{-2\mu\gamma} \mathbb{P}_{-\mu}(S_{m(\gamma)} \leq -\gamma)p(\mu) d\mu + \int_{-\infty}^{-\delta} e^{2\mu\gamma} \mathbb{P}_{-\mu}(S_{m(\gamma)} \geq \gamma)p(\mu) d\mu \\ &\leq \int_\delta^\infty e^{-2\mu\gamma} p(\mu) d\mu + \int_{-\infty}^{-\delta} e^{2\mu\gamma} p(\mu) d\mu \leq e^{-2\delta\gamma} \int_{|\mu|>\delta} p(\mu) d\mu = o(\gamma^{-2}). \end{aligned}$$
■

Lemma 2. As $\gamma \rightarrow \infty$,

$$\int_{|\mu|\leq\delta} |\mathbb{E}_\mu[b(\gamma)]\mu| d\mu = O(\gamma^{-2}).$$

Proof. Notice that, using Proposition 1,

$$\begin{aligned} \int_{|\mu|\leq\delta} |\mathbb{E}_\mu[b(\gamma)]\mu| d\mu &\leq \int_{-\delta}^0 |\mathbb{E}_\mu[b_+(\gamma)]\mu| d\mu + \int_0^\delta |\mathbb{E}_\mu[b_-(\gamma)]\mu| d\mu \\ &\leq \int_{-\delta}^0 \left| e^{2\mu\gamma} \mathbb{P}_{-\mu}(S_{m(\gamma)} \geq \gamma) \mu \right| d\mu + \int_0^\delta \left| e^{-2\mu\gamma} \mathbb{P}_{-\mu}(S_{m(\gamma)} \leq -\gamma) \mu \right| d\mu \\ &\leq \int_{-\delta}^0 \left| e^{2\mu\gamma} \mu \right| d\mu + \int_0^\delta \left| e^{-2\mu\gamma} \mu \right| d\mu = \frac{1}{2\gamma^2} - \frac{e^{-2\delta\gamma}}{2\gamma^2} - \frac{\delta e^{-2\delta\gamma}}{\gamma}. \end{aligned}$$

The result follows. ■

A.3. Expected Number of Inner Samples

The asymptotic characterization of the number of inner samples provided by Theorem 3 is that, as $\gamma \rightarrow \infty$, $\bar{m}(\gamma) = O(\gamma \log \gamma)$. We will establish this via an analysis of (28). In particular, we have that

$$(34) \quad \bar{m}(\gamma) = \int_{|\mu|>\gamma^{-1}} \mathbb{E}_\mu[m(\gamma)]p(\mu) d\mu + \int_{|\mu|\leq\gamma^{-1}} \mathbb{E}_\mu[m(\gamma)]p(\mu) d\mu.$$

Theorem 3 will follow by applying Lemmas 4 and 5, established below, to (34).

To this end, the following result will be helpful.

Lemma 3. *Suppose Y_1, Y_2, \dots are i.i.d. random variables under the probability measure \mathbb{P}_μ , with $\mathbb{E}_\mu[Y_1] = \mu$ and $\mathbb{E}_\mu[Y_1^2] < \infty$. Define, for $m \geq 0$, the partial sum $S_m \triangleq \sum_{j=1}^m Y_j$, and, for $\gamma > 0$, the one-sided hitting times*

$$m_+(\gamma) \triangleq \inf \{m > 0 : S_m > \gamma\}, \quad m_\pm(\gamma) \triangleq \inf \{m > 0 : |S_m| > \gamma\}.$$

(i) (Lorden, 1970) *Suppose that $\mu > 0$. Then, if $x^+ \triangleq \max(x, 0)$,*

$$\sup_{\gamma > 0} \mathbb{E}_\mu[S_{m_+(\gamma)} - \gamma] \leq \frac{\mathbb{E}_\mu[(Y_1^+)^2]}{\mu}.$$

(ii) (Pruitt, 1981) *There exist constants V_1 and V_2 (independent of the distribution of Y_1) such that, if $K_\mu(\gamma) \triangleq \gamma^{-2} \mathbb{E}_\mu[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma\}}]$,*

$$\mathbb{E}_\mu[m_\pm(\gamma)] \leq \frac{V_1}{K_\mu(\gamma)}, \quad \mathbb{P}_\mu \left(\max_{1 \leq m \leq n} |S_m| \leq \gamma \right) \leq \frac{V_2}{(nK_\mu(\gamma))^3}.$$

(iii) (Gut, 1974)

$$\mathbb{E} \left[\left(Y_{m_+(\gamma)}^+ \right)^2 \right] \leq \mathbb{E}[m_+(\gamma)] \mathbb{E} \left[\left(Y_1^+ \right)^2 \right].$$

Lemma 4. *As $\gamma \rightarrow \infty$,*

$$\int_{|\mu| > \gamma^{-1}} \mathbb{E}_\mu[m(\gamma)] p(\mu) d\mu = O(\gamma \log \gamma).$$

Proof. Note that, since Y_1 has mean μ and unit variance under the distribution \mathbb{P}_μ ,

$$(35) \quad \frac{\mathbb{E}_\mu[(Y_1^+)^2]}{\mu} \leq \frac{\mathbb{E}_\mu[|Y_1|^2]}{\mu} = \frac{1 + \mu^2}{\mu}.$$

Further, define the one-sided hitting times $m_+(\gamma)$ and $m_-(\gamma)$ as in Lemma 3. By the optional stopping theorem,

$$\mathbb{E}_\mu[S_{m_+(\gamma)}] = \mu \mathbb{E}_\mu[m_+(\gamma)], \quad \text{if } \mu > 0; \quad \mathbb{E}_\mu[S_{m_-(\gamma)}] = \mu \mathbb{E}_\mu[m_-(\gamma)], \quad \text{if } \mu < 0.$$

Then, since $m(\gamma) \leq m_+(\gamma)$ and $m(\gamma) \leq m_-(\gamma)$, we have that

$$\begin{aligned}
\int_{|\mu|>\gamma^{-1}} \mathbf{E}_\mu[m(\gamma)]p(\mu) d\mu &\leq \int_{-\infty}^{-\gamma^{-1}} \mathbf{E}_\mu[m_-(\gamma)]p(\mu) d\mu + \int_{\gamma^{-1}}^{\infty} \mathbf{E}_\mu[m_+(\gamma)]p(\mu) d\mu \\
&= \int_{-\infty}^{-\gamma^{-1}} \frac{\mathbf{E}_\mu[S_{m_-(\gamma)}]}{\mu} p(\mu) d\mu + \int_{\gamma^{-1}}^{\infty} \frac{\mathbf{E}_\mu[S_{m_+(\gamma)}]}{\mu} p(\mu) d\mu \\
(36) \quad &\leq \int_{-\infty}^{-\gamma^{-1}} \left(\frac{1+\mu^2}{\mu^2} - \frac{\gamma}{\mu} \right) p(\mu) d\mu + \int_{\gamma^{-1}}^{\infty} \left(\frac{1+\mu^2}{\mu^2} + \frac{\gamma}{\mu} \right) p(\mu) d\mu \\
&= \int_{|\mu|>\gamma^{-1}} \left(1 + \frac{1}{\mu^2} + \frac{\gamma}{|\mu|} \right) p(\mu) d\mu.
\end{aligned}$$

Here, the final inequality follows from (35) and Part (i) of Lemma 3.

Now, without loss of generality, assume that $\gamma > \delta^{-1}$. Recalling U_0 from (30), we have that

$$\begin{aligned}
&\int_{|\mu|>\gamma^{-1}} \mathbf{E}_\mu[m(\gamma)]p(\mu) d\mu \\
(37) \quad &\leq \int_{\delta \geq |\mu|>\gamma^{-1}} \left(1 + \frac{1}{\mu^2} + \frac{\gamma}{|\mu|} \right) p(\mu) d\mu + \int_{|\mu|>\delta} \left(1 + \frac{1}{\mu^2} + \frac{\gamma}{|\mu|} \right) p(\mu) d\mu \\
&\leq 2U_0 \int_{\gamma^{-1}}^{\delta} \left(1 + \frac{1}{\mu^2} + \frac{\gamma}{\mu} \right) d\mu + \left(1 + \frac{1}{\delta^2} + \frac{\gamma}{\delta} \right) \int_{|\mu|>\delta} p(\mu) d\mu \\
&= 2U_0 \left(\delta - \gamma^{-1} + \gamma - \delta^{-1} + \gamma \log \delta + \gamma \log \gamma \right) + 1 + \delta^{-2} + \gamma \delta^{-1} = O(\gamma \log \gamma).
\end{aligned}$$

■

Lemma 5. As $\gamma \rightarrow \infty$,

$$\int_{|\mu| \leq \gamma^{-1}} \mathbf{E}_\mu[m(\gamma)]p(\mu) d\mu = O(\gamma).$$

Proof. Here, we will apply Part (ii) of Lemma 3. Without loss of generality, assume that $\gamma > 1$.

Then, $|\mu| < 1$ in the region of integration and thus $\gamma - |\mu| > 0$. $K_\mu(\gamma)$ from Lemma 3 satisfies

$$\begin{aligned}
(38) \quad K_\mu(\gamma) &= \gamma^{-2} \mathbf{E}_\mu \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma\}} \right] = \gamma^{-2} \mathbf{E}_0 \left[|Y_1 + \mu|^2 \mathbb{I}_{\{|Y_1 + \mu| \leq \gamma\}} \right] \geq \gamma^{-2} \mathbf{E}_0 \left[|Y_1 + \mu|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}} \right] \\
&\geq \gamma^{-2} \left(\mathbf{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}} \right] + 2\mu \mathbf{E}_0 \left[Y_1 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}} \right] \right) = \gamma^{-2} \mathbf{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}} \right].
\end{aligned}$$

Here, we have used the fact that under \mathbf{P}_0 , $Y_1 \sim N(0, 1)$, hence $\mathbf{E}_0[Y_1 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}}] = 0$.

Then, from Part (ii) of Lemma 3, since $m(\gamma) \leq m_\pm(\gamma)$,

$$\mathbf{E}_\mu[m(\gamma)] \leq \mathbf{E}_\mu[m_\pm(\gamma)] \leq \frac{V_1 \gamma^2}{\mathbf{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}} \right]}.$$

Without loss of generality, assume that $\gamma > \delta^{-1}$, and recall U_0 from (30). Then,

$$(39) \quad \begin{aligned} \int_{|\mu| \leq \gamma^{-1}} \mathbb{E}_\mu[m(\gamma)] p(\mu) d\mu &\leq \int_{|\mu| \leq \gamma^{-1}} \frac{V_1 \gamma^2}{\mathbb{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}} \right]} p(\mu) d\mu \\ &\leq \frac{V_1 \gamma^2}{\mathbb{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - \delta\}} \right]} \int_{|\mu| \leq \gamma^{-1}} p(\mu) d\mu \leq \frac{2U_0 V_1 \gamma}{\mathbb{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - \delta\}} \right]}. \end{aligned}$$

Notice that $\gamma > 1 > \delta$ is assumed before. By the monotone convergence theorem,

$$(40) \quad \lim_{\gamma \rightarrow \infty} \mathbb{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - \delta\}} \right] = \mathbb{E}_0 \left[|Y_1|^2 \right] = 1.$$

The result follows. ■

A.4. Realized Number of Inner Samples

In this section, we will establish Theorem 4, which provides a probabilistic bound on the realized number of inner stage samples per scenario. Our proof relies on the following lemma, which bounds the second moment of the number of inner stage samples per scenario.

Lemma 6. *As $\gamma \rightarrow \infty$, $\mathbb{E} [m(\gamma)^2] = O(\gamma^3)$.*

We will defer the proof of Lemma 6 for the moment, and first employ this lemma to prove Theorem 4.

Proof of Theorem 4. Fix $\epsilon > 0$ and suppose that $\gamma \geq \gamma_0$. Then, by Chebyshev's inequality,

$$\begin{aligned} \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n m_i \geq (C_0 + \epsilon) \gamma \log \gamma \right) &\leq \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n (m_i - \bar{m}(\gamma)) \right| \geq \epsilon \gamma \log \gamma \right) \\ &\leq \frac{\text{Var} (m(\gamma))}{n (\epsilon \gamma \log \gamma)^2} \leq \frac{\mathbb{E} [m(\gamma)^2]}{n (\epsilon \gamma \log \gamma)^2}. \end{aligned}$$

By Lemma 6, there exist constants $C'_0, \gamma'_0 > 0$ so that if $\gamma \geq \gamma'_0$, $\mathbb{E} [m(\gamma)^2] \leq C'_0 \gamma^3$. Then, if $\gamma \geq \max\{\gamma_0, \gamma'_0, \gamma_1\}$, we have that

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n m_i \geq (C_0 + \epsilon) \gamma \log \gamma \right) \leq \frac{C'_0}{C_1 (\epsilon \log \gamma)^2},$$

which can be made arbitrarily small with sufficiently large γ . ■

To prove Lemma 6, consider the decomposition

$$(41) \quad \mathbb{E} [m(\gamma)^2] = \int_{|\mu| > \gamma^{-1}} \mathbb{E}_\mu [m(\gamma)^2] p(\mu) d\mu + \int_{|\mu| \leq \gamma^{-1}} \mathbb{E}_\mu [m(\gamma)^2] p(\mu) d\mu.$$

Lemma 6 will follow by applying Lemmas 7 and 8, established below, to (41).

Lemma 7. As $\gamma \rightarrow \infty$,

$$\int_{|\mu| > \gamma^{-1}} \mathbf{E}_\mu \left[m(\gamma)^2 \right] p(\mu) d\mu = O(\gamma^3).$$

Proof. We proceed as in the proof of Lemma 4. Using the stopping times $m_+(\gamma)$ and $m_-(\gamma)$ defined there, we have

$$\int_{|\mu| > \gamma^{-1}} \mathbf{E}_\mu \left[m(\gamma)^2 \right] p(\mu) d\mu \leq \int_{-\infty}^{-\gamma^{-1}} \mathbf{E}_\mu \left[m_-(\gamma)^2 \right] p(\mu) d\mu + \int_{\gamma^{-1}}^{\infty} \mathbf{E}_\mu \left[m_+(\gamma)^2 \right] p(\mu) d\mu.$$

First, consider the case when $\mu > 0$. By the optional stopping theorem applied to the quadratic martingale $(S_m - \mu m)^2 - m$, we have that $\mathbf{E}_\mu [(S_{m_+(\gamma)} - \mu m_+(\gamma))^2] = \mathbf{E}_\mu [m_+(\gamma)]$. Now, for any real numbers $a, b \in \mathbb{R}$, we have that $(a + b)^2 \leq 2(a^2 + b^2)$. Therefore,

$$\mathbf{E}_\mu \left[m_+(\gamma)^2 \right] \leq \frac{2}{\mu^2} \left(\mathbf{E}_\mu \left[(S_{m_+(\gamma)} - \mu m_+(\gamma))^2 \right] + \mathbf{E}_\mu \left[S_{m_+(\gamma)}^2 \right] \right) = \frac{2}{\mu^2} \left(\mathbf{E}_\mu [m_+(\gamma)] + \mathbf{E}_\mu \left[S_{m_+(\gamma)}^2 \right] \right)$$

Using the fact that $S_{m_+(\gamma)} \leq \gamma + Y_{m_+(\gamma)}^+$, Part (iii) of Lemma 3, and (35),

$$\begin{aligned} \mathbf{E}_\mu \left[m_+(\gamma)^2 \right] &\leq \frac{2}{\mu^2} \left(\mathbf{E}_\mu [m_+(\gamma)] + \mathbf{E}_\mu \left[(\gamma + Y_{m_+(\gamma)}^+)^2 \right] \right) \\ &\leq \frac{2}{\mu^2} \left(\mathbf{E}_\mu [m_+(\gamma)] + 2\gamma^2 + 2\mathbf{E}_\mu \left[(Y_{m_+(\gamma)}^+)^2 \right] \right) \\ &\leq \frac{2}{\mu^2} \left(\mathbf{E}_\mu [m_+(\gamma)] + 2\gamma^2 + 2\mathbf{E}_\mu [m_+(\gamma)] \mathbf{E}_\mu \left[(Y_1^+)^2 \right] \right) \\ &\leq \frac{2}{\mu^2} \left(\mathbf{E}_\mu [m_+(\gamma)] + 2\gamma^2 + 2(\mu^2 + 1)\mathbf{E}_\mu [m_+(\gamma)] \right) = \left(\frac{6}{\mu^2} + 4 \right) \mathbf{E}_\mu [m_+(\gamma)] + \frac{4\gamma^2}{\mu^2}. \end{aligned}$$

By similar consideration of the symmetric case where $\mu < 0$, we have, repeating the calculation in (36),

$$\int_{|\mu| > \gamma^{-1}} \mathbf{E}_\mu \left[m(\gamma)^2 \right] p(\mu) d\mu \leq \int_{|\mu| > \gamma^{-1}} \left[\left(\frac{6}{\mu^2} + 4 \right) \left(1 + \frac{1}{\mu^2} + \frac{\gamma}{|\mu|} \right) + \frac{4\gamma^2}{\mu^2} \right] p(\mu) d\mu.$$

Without loss of generality, assume that $\gamma > \delta^{-1}$. Then, as in (37),

$$\begin{aligned} \int_{|\mu| > \gamma^{-1}} \mathbf{E}_\mu \left[m(\gamma)^2 \right] p(\mu) d\mu &\leq 2U_0 \int_{\gamma^{-1}}^{\delta} \left[\left(\frac{6}{\mu^2} + 4 \right) \left(1 + \frac{1}{\mu^2} + \frac{\gamma}{\mu} \right) + \frac{4\gamma^2}{\mu^2} \right] d\mu \\ &\quad + \left(\frac{6}{\delta^2} + 4 \right) \left(1 + \frac{1}{\delta^2} + \frac{\gamma}{\delta} \right) + \frac{4\gamma^2}{\delta^2} = O(\gamma^3). \end{aligned}$$

■

Lemma 8. As $\gamma \rightarrow \infty$,

$$\int_{|\mu| \leq \gamma^{-1}} \mathbf{E}_\mu \left[m(\gamma)^2 \right] p(\mu) d\mu = O(\gamma^3).$$

Proof. We proceed as in Lemma 5. Without loss of generality, assume that $\gamma > 1$. Observe that

$m(\gamma) \leq m_{\pm}(\gamma)$, since the latter is an exit time for a larger set than the former. Then, using summation by parts,

$$\begin{aligned} \mathbf{E}_{\mu} \left[m(\gamma)^2 \right] &\leq \mathbf{E}_{\mu} \left[m_{\pm}(\gamma)^2 \right] = \sum_{n=1}^{\infty} n^2 \mathbf{P}_{\mu}(m_{\pm}(\gamma) = n) \\ &= 1 + \sum_{n=1}^{\infty} (2n+1) \mathbf{P}_{\mu}(m_{\pm}(\gamma) > n) \\ &= 1 + \sum_{n=1}^{\infty} (2n+1) \mathbf{P}_{\mu} \left(\max_{1 \leq m \leq n} |S_m| \leq \gamma \right). \end{aligned}$$

Using Part (ii) of Lemma 3, for any integer $N \geq 1$,

$$\begin{aligned} \mathbf{E}_{\mu} \left[m(\gamma)^2 \right] &\leq \sum_{n=0}^{N-1} (2n+1) + \frac{V_2}{K_{\mu}(\gamma)^3} \sum_{n=N}^{\infty} \frac{2n+1}{n^3} \\ &\leq N^2 + \frac{3V_2}{K_{\mu}(\gamma)^3} \sum_{n=N}^{\infty} \frac{1}{n^2} \leq N^2 + \frac{3V_2}{K_{\mu}(\gamma)^3(N-1)}. \end{aligned}$$

Since $K_{\mu}(\gamma) \leq 1$, we may take $N \triangleq \lceil 3/K_{\mu}(\gamma) \rceil$, so that $N-1 \geq 1/K_{\mu}(\gamma)$. Then, there exists a constant W_0 so that

$$\mathbf{E}_{\mu} \left[m(\gamma)^2 \right] \leq \frac{W_0}{K_{\mu}(\gamma)^2} \leq \frac{W_0 \gamma^4}{\left(\mathbf{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}} \right] \right)^2},$$

using (38).

Without loss of generality, assume that $\gamma > \delta^{-1}$. Then, as in (39),

$$\int_{|\mu| \leq \gamma^{-1}} \mathbf{E}_{\mu} \left[m(\gamma)^2 \right] p(\mu) d\mu \leq \int_{|\mu| \leq \gamma^{-1}} \frac{W_0 \gamma^4}{\left(\mathbf{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - |\mu|\}} \right] \right)^2} p(\mu) d\mu \leq \frac{2U_0 W_0 \gamma^3}{\left(\mathbf{E}_0 \left[|Y_1|^2 \mathbb{I}_{\{|Y_1| \leq \gamma - \delta\}} \right] \right)^2}.$$

The result follows from (40). ■