2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

Optimization of Utility-based Shortfall Risk: A Non-asymptotic Viewpoint

Sumedh Gupte[1] and Prashanth L. A.[2] and Sanjay P. Bhat[3]

***Abstract*— We consider the problems of estimation and optimization of utility-based shortfall risk (UBSR), which is a popular risk measure in finance. In the context of UBSR estimation, we derive a non-asymptotic bound on the mean-squared error of the classical sample average approximation (SAA) of UBSR. Next, in the context of UBSR optimization, we derive an expression for the UBSR gradient under a smooth parameterization. This expression is a ratio of expectations, both of which involve the UBSR. We use SAA for the numerator as well as denominator in the UBSR gradient expression to arrive at a biased gradient estimator. We derive non-asymptotic bounds on the estimation error, which show that our gradient estimator is asymptotically unbiased. We incorporate the aforementioned gradient estimator into a stochastic gradient (SG) algorithm for UBSR optimization. Finally, we derive non-asymptotic bounds that quantify the rate of convergence of our SG algorithm for UBSR optimization.**

***Index Terms*— Utility-based shortfall risk, risk estimation, stochastic optimization, non-asymptotic bounds, stochastic gradient.**

I. Introduction

Optimizing risk is important in several application domains, e.g., finance, transportation to name a few. Financial applications rely heavily on efficient risk assessment techniques, and employ multitude of risk measures for risk estimation. Risk optimization involves risk estimation as a sub-procedure for finding solutions to optimal decision-making problems in finance. Value-at-Risk (VaR) [1], Conditional Value-at-Risk (CVaR) [2] are two popular risk measures. The risk measure VaR, which is a quantile of the underlying distribution, is not the preferred choice owing to the fact that it is not sub-additive. In a financial context, the latter property implies that diversification does not increase risk. CVaR as a risk measure satisfies sub-additivity property and falls in the category of coherent risk measures [3].

A class of risk measures that subsumes coherency is convex risk measures [4]. A prominent convex risk measure is utility-based shortfall risk (UBSR). In this paper, we consider the problems of UBSR estimation and optimization. UBSR is a convex risk measure that has a few advantages over the popular CVaR risk measure, namely (i) UBSR is invariant under randomization, while CVaR is not, see [5]; (ii) Unlike CVaR, which only considers the values that the underlying random variable takes beyond VaR, the loss function in UBSR can be chosen to encode the risk preference for each value that the underlying random variable takes. Thus, in the context of both risk estimation and optimization, UBSR is a viable alternative to the industry standard risk measures, namely, VaR and CVaR.

The existing works on UBSR are restricted to the case where the underlying random variables are bounded, cf. [6]. In this paper, we extend the UBSR formalism to unbounded random variables. We now summarize our contributions below.

1) We extend UBSR to cover unbounded random variables that satisfy certain integrability requirements, and establish conditions under which UBSR is a convex risk measure.

2) For a sample average approximation (SAA) of UBSR, which is proposed earlier in the literature, we derive a mean-squared error bound under a weaker assumption on the underlying loss function. More precisely, we assume the loss function to be bounded above by a quadratic function, while the corresponding bound in [7] assumed linear growth.

3) For the problem of UBSR optimization with a vector parameter, we derive an expression for the gradient of UBSR. Using this expression, we propose an m-sample gradient estimator. We establish $\mathcal{O}(1/m)$ and $\mathcal{O}(1/\sqrt{m})$ mean-squared error bounds for Lipschitz and smooth loss functions, respectively.

4) We design a stochastic gradient (SG) algorithm using the gradient estimator above, and derive a non-asymptotic bound of $\mathcal{O}(1/n)$ under a strong convexity assumption with a Lipschitz loss function. Here n denotes the number of iterations of the SG algorithm for UBSR optimization. A similar proof leads to a $\mathcal{O}(1/\sqrt{n})$ bound for a smooth loss function.

Related work. The authors in [8] introduced UBSR for bounded random variables and the authors in [9] illustrated several desirable properties of UBSR. In real-world financial markets, the financial positions continuously evolve over time, and so must their risk estimates. In [10], the authors showed that UBSR can be used for dynamic evaluation of such financial positions. In [5], the authors have proposed estimators based on a stochastic root finding procedure and they provide only asymptotic convergence guarantees. In [11], the authors have used sample average approximation (SAA) procedure for UBSR estimation and have proposed an estimator for the UBSR derivative which can be used for risk optimization in the case of a scalar decision parameter. They establish asymptotic convergence guarantees for UBSR estimation and also show that the UBSR derivative is asymptotically unbiased. In [6], the authors perform non-asymptotic analysis for the scalar UBSR optimization, while employing a stochastic root finding technique for UBSR

[1] TCS Research, Hyderabad, India, sumedh.gupte@tcs.com
[2] Department of Computer Science and Engineering, IIT Madras, India, prashla@cse.iitm.ac.in
[2] TCS Research, Hyderabad, India, sanjay.bhat@tcs.com

estimation. In comparison to these works, we would like to the note the following aspects: (i) Unlike [5], [11], we provide non-asymptotic bounds on the mean-squared error of the UBSR estimate from a procedure that is computationally efficient; (ii) We consider UBSR optimization for a vector parameter, while earlier works (cf. [11], [6]) consider the scalar case; (iii) We analyze a SG-based algorithm in the non-asymptotic regime for UBSR optimization, while [11] provides an asymptotic guarantee for the UBSR derivative estimate; (iv) In [6], UBSR optimization using a gradient-based algorithm has been proposed for the case of scalar parameterization. Unlike [6], we derive a general (multi-variate) expression for the UBSR gradient, leading to an estimator that is subsequently employed in the stochastic gradient algorithm mentioned above. A vector parameter makes the bias/variance analysis of UBSR gradient estimate challenging as compared to the scalar counterpart, which is analyzed in [6].

The rest of the paper is organized as follows: In Section II, we introduce the notations. In Section III we extend the UBSR risk measure for unbounded random variables. In Section IV, we describe the SAA-based UBSR estimation technique, and derive error bounds on the estimator. In Section V, we derive the UBSR gradient theorem, and propose as well as analyze a stochastic gradient scheme for UBSR minimization. In Section VI, we provide the concluding remarks. Due to space constraints, we provide detailed proofs in the longer version of this paper, available in [12].

II. Preliminaries

For $p \in [1, \infty)$, the p-norm of a vector $\mathbf{v} \in \mathbb{R}^d$ is given by

$$\left\|\mathbf{v}\right\|_p \triangleq \left(\sum_{i=1}^{d} |v_i|^p\right)^{\frac{1}{p}},$$

while $\left\|\mathbf{v}\right\|_\infty$ denotes the supremum norm. Matrix norms induced by the vector p-norm are denoted by $\left\|\cdot\right\|_p$, and the special cases of $p = 1$ and $p = \infty$ denote the maximum absolute column sum and maximum absolute row sum, respectively. The spectral norm is denoted by $\left\|\cdot\right\|$.

Let $(\Omega, \mathcal{F}, P)$ be a standard Borel probability space. Let L^0 denote the space of $\mathcal{F}$-measurable, real random variables and let $\mathbb{E}(\cdot)$ denote the expectation under P. For $p \in [1, \infty)$, let $\left(L^p, \left\|\cdot\right\|_{L^p})\right)$ denote the normed vector space of random variables $X : \Omega \to \mathbb{R}$ in L^0 for which $\left\|X\right\|_{L^p} \triangleq \left(\mathbb{E}\left[|X|^p\right]\right)^{\frac{1}{p}}$ is finite. Further, we let $\left(L^\infty, \left\|\cdot\right\|_{L^\infty}\right)$ denote the normed vector space of random variables $X : \Omega \to \mathbb{R}$ in L^0, for which, $\left\|X\right\|_{L^\infty} \triangleq \inf\{M \in \mathbb{R} : |X| \leq M \text{ a.s.}\}$ is finite. Let $p \in [1, \infty)$ and let $\mathbf{Z}$ be a random vector such that each Z_i is $\mathcal{F}$-measurable and has finite p^{th} moment. Then the L^p-norm of $\mathbf{Z}$ is defined by $\left\|\mathbf{Z}\right\|_{L^p} \triangleq \left(\mathbb{E}\left[\,\left\|\mathbf{Z}\right\|_p^p\right]\right)^{\frac{1}{p}}$.

Let μ_X and μ_Y denote the marginal distributions of random vectors $\mathbf{X}$ and $\mathbf{Y}$ respectively. Let $\mathcal{H}(\mu_\mathbf{X}, \mu_\mathbf{Y})$ denote the set of all joint distributions having $\mu_\mathbf{X}$ and $\mu_\mathbf{Y}$ as the marginals. Then, for every $p \geq 1$, $\mathcal{W}_p(\mu_\mathbf{X}, \mu_\mathbf{Y}) \triangleq \left(\inf\left\{\int \|x - y\|^p\, \eta(dx, dy) : \eta \in \mathcal{H}\right\}\right)^{1/p}$ denotes the p-Wasserstein distance associated with $\mathbf{X}$ and $\mathbf{Y}$. Note that vectors and random vectors are distinguished from their scalar counterparts by the use of boldface fonts.

Throughout this paper, we shall use $l : \mathbb{R} \to \mathbb{R}$ to denote a continuous function that models a loss function, for instance, of a financial position. Let $\mathcal{X}_l$ denote the space of random variables $X \in L^0$ for which the collection of random variables $\{l(-X - t) : t \in \mathbb{R}\}$ is uniformly integrable. With P being finite, uniform integrability implies that for every $X \in \mathcal{X}_l$, we have: $\sup_{t\in\mathbb{R}} \int_\Omega |l(-X - t)|dP < \infty$. While the integrability condition is not necessary for defining UBSR, we have incorporated it into the definition of $\mathcal{X}_l$. This condition is used to characterize the UBSR as a unique root of a decreasing function, and this characterization is useful in the analysis of UBSR estimation and optimization techniques. Note that $L^\infty \subseteq \mathcal{X}_l$ holds for any continuous function l. Under suitable assumptions on l, $\mathcal{X}_l$ also contains unbounded random variables. For example, choosing a Lipschitz continuous loss function l ensures that $\mathcal{X}_l$ contains all square-integrable random variables.

III. UBSR for Unbounded Random Variables

Let $X \in \mathcal{X}_l$ denote the random variable representing the payoff or utility, e.g., of a financial position. Formally, a financial position is a mapping $X : \Omega \to \mathbb{R}$, where $X(\omega)$ is the net worth of the position after its realisation. Then, $-X$ denotes the loss incurred, and the UBSR is the smallest amount to be added to X so that the expected loss is below a certain acceptable threshold, say λ. We formalize the notion of UBSR [4] below.

Definition 1: With the loss function l and a given threshold $\lambda \in \mathbb{R}$, the risk measure UBSR of $X \in \mathcal{X}_l$ is defined as the function $SR_{l,\lambda} : \mathcal{X}_l \to \mathbb{R}$, where

$$SR_{l,\lambda}(X) \triangleq \inf\{t \in \mathbb{R} | \mathbb{E}[l(-X - t)] \leq \lambda\}. \tag{1}$$

Following [13], we define the *acceptance set* associated with the UBSR risk measure as follows:

$$\mathcal{A}_{l,\lambda} = \{X \in \mathcal{X}_l : SR_{l,\lambda}(X) \leq 0\}. \tag{2}$$

Observe that the set $\mathcal{A}_{l,\lambda}$ contains all random variables X whose expected loss $\mathbb{E}[l(-X)]$ does not exceed λ.

As an example, with $l(x) = \exp(\beta x)$ and $\lambda = 1$, $SR_{l,\lambda}(X)$ is identical to the entropic risk measure [14], which is a coherent risk measure and enjoys several advantages over the standard risk measures VaR and CVaR. The other popular choice is $l(x) = \frac{x^p}{p}$ for $x \geq 0$ and 0 otherwise, with $p > 1$.

A. Characterization of UBSR

We discuss the problem of quantifying the risk $SR_{l,\lambda}(X)$ of a financial position $X \in \mathcal{X}_l$. To this end, we recall an alternative characterization for $SR_{l,\lambda}(X)$. Consider a real-valued function $g_X : \mathbb{R} \to \mathbb{R}$ associated with the random variable X as

$$g_X(t) \triangleq \mathbb{E}[l(-X - t)] - \lambda. \tag{3}$$

Note that $SR_{l,\lambda}(X)$, if it exists, is a root of $g_X(\cdot)$. We now introduce the following assumptions on X, l and λ under which the existence and uniqueness of this root follows.

Assumption 1: There exists $t_{\mathrm{u}}, t_{\mathrm{l}} \in \mathbb{R}$ such that $g_X(t_{\mathrm{u}}) \leq 0$ and $g_X(t_{\mathrm{l}}) \geq 0$.

Assumption 2: The function l is an increasing function such that $\forall t \in \mathbb{R}, P(l'(-X-t) > 0) > 0$.

Assumption 3: The function l is continuously differentiable, and the collection of random variables $\{l'(-X-t) : t \in R\}$ is uniformly integrable.

The first assumption is required for the existence of UBSR, while the other assumptions ensures that g_X is strictly decreasing, a property we exploit in UBSR estimation. The last assumption is required for interchanging the derivative and the expectation. Similar assumptions have been made in the context of UBSR estimation in [11], [6] that also assume that random variables have bounded support. In contrast, we extend the analysis to unbounded random variables using an added assumption of uniform integrability.

The result below establishes that UBSR is the unique root of the function g_X defined in eq. (3).

Proposition 3.1: Let Assumptions 1 to 3 hold. Then g_X is continuous and strictly decreasing, and the unique root of g_X coincides with $SR_{l,\lambda}(X)$.

See [12, Appendix I] for the proof.

Remark 1: A similar result for bounded random variables has been stated in [8, Proposition 4.104] under the assumption that the loss function l is convex and strictly increasing. We generalize this to unbounded random variables. Unlike [8], our proof does not require the convexity assumption, and we relax the strictly increasing assumption by replacing it with Assumption 2.

B. Wasserstein Distance Bounds on UBSR

We provide results for UBSR estimation and optimization under two different assumptions on the underlying loss functions. These assumptions are specified below.

Assumption 4: l is L_1-Lipschitz, i.e., there exists $L_1 > 0$ such that for every $x, y \in \mathbb{R},\ |l(x) - l(y)| \leq L_1 |x - y|$.

Assumption 4′: l is convex and L_2-smooth, and has sub-linear derivative, i.e., there exist $L_2 > 0, a > 0$ and $b > 0$ such that for every $x, y \in \mathbb{R}$, we have

$$l(y) - l(x) - l'(x)(y-x) \leq \frac{L_2}{2}(y-x)^2, \qquad (4)$$

$$b \leq l'(x) \leq a|x| + b, \qquad (5)$$

where b simultaneously satisfies Assumption 5.

Using the assumptions above, we derive bounds on the difference between the UBSR values of two random variables X and Y. The bounds obtained are in terms of the 1-Wasserstein and 2-Wasserstein distances between the corresponding marginal distributions μ_X, μ_Y. Under Assumption 4, such a result has been obtained in [7]. We provide a simpler proof which generalizes to both, Lipschitz and non-Lipschitz loss functions. The Lipschitzness assumption implies that the derivative is bounded, which is often restrictive. Assumption 4′ relaxes this condition to allow for unbounded, but sub-linear derivatives.

In addition, we require the following assumption for all the results.

Assumption 5: There exists $b > 0$ such that for every $x, y \in \mathbb{R}, y \geq x \implies l(y) - l(x) \geq b(y-x)$.

The result below provides a Wasserstein distance bound for smooth loss functions. This bound will be used in Section IV for establishing error bounds for UBSR estimation.

Lemma 1: Suppose Assumption 4′ and Assumption 5 hold, and there exists $T > 0$ such that $\|X\|_{L^2} \leq T$ for every $X \in \mathcal{X}_l$. Then for every $X, Y \in \mathcal{X}_l$, we have

$$|SR_{l,\lambda}(X) - SR_{l,\lambda}(Y)| \leq \frac{L_2}{2b}\mathcal{W}_2^2(\mu_X, \mu_Y) + \frac{aT}{b}\mathcal{W}_2(\mu_X, \mu_Y), \qquad (6)$$

where a, b, L_2 are as specified in assumptions 4 and 5. See [12, Appendix I] for the proof.

In [7], the authors establish the following bound under assumptions 4 and 5:

$$|SR_{l,\lambda}(X) - SR_{l,\lambda}(Y)| \leq \frac{L_1}{b}\mathcal{W}_1(\mu_X, \mu_Y). \qquad (7)$$

While the inequality above is useful for deriving bounds for UBSR estimation, Assumption 4 is restrictive since it excludes quadratic losses, which appear naturally in mean-variance optimization. Assumption 4′ covers such loss functions, and we bound UBSR difference between two distributions using the 2-Wasserstein distance in this case.

Next we briefly discuss the properties that are associated with risk measures. These properties satisfy certain desirable, investor preferences. Readers are referred to [8], [13] for a detailed study.

C. Convexity of the UBSR Risk Measure

We define the notions of monetary and convex risk measures below [8]. To do so, we use $\mathcal{X}$ to denote an arbitrary set of random variables.

Definition 2: A mapping $\rho : \mathcal{X} \to \mathbb{R}$ is called a monetary measure of risk if it satisfies the following two conditions.

1) Monotonicity: For all $X_1, X_2 \in \mathcal{X}$ such that $X_1 \leq X_2$ a.s., we have $\rho(X_1) \geq \rho(X_2)$.
2) Cash invariance: For all $X \in \mathcal{X}$ and $m \in \mathbb{R}$, we have $\rho(X+m) = \rho(X) - m$.

Definition 3: A monetary risk measure ρ is convex if for every $X_1, X_2 \in \mathcal{X}$, and $\alpha \in [0,1]$, the following holds:

$$\rho(\alpha X_1 + (1-\alpha)X_2) \leq \alpha\rho(X_1) + (1-\alpha)\rho(X_2). \qquad (8)$$

In this paper, we focus on UBSR, which is a convex risk measure. Existing works [8] have shown that UBSR risk measure is convex for the restricted case of bounded random variables ($\mathcal{X} \subseteq L^\infty$). Using a novel proof technique in the following proposition, we extend this convexity of UBSR to unbounded random variables.

Proposition 3.2: Suppose Assumptions 1 to 3 hold for l and every $X \in \mathcal{X}_l$. Then $SR_{l,\lambda}(\cdot)$ is a monetary risk measure. In addition, if l is also convex, then $\mathcal{A}_{l,\lambda}$ is convex, and $SR_{l,\lambda}(\cdot)$ is a convex risk measure.

See [12, Appendix I] for the proof. This result is useful in showing that the problem of UBSR optimization falls under the class of stochastic convex optimization problems.

IV. UBSR Estimation

In this section, we discuss techniques to compute UBSR for a given random variable X. In practice, the true distribution of X is unavailable, and instead one relies on the samples of X to estimate the UBSR. We use the sample average approximation (SAA) technique [15] for UBSR estimation of a random variable $X \in \mathcal{X}_l$. Such a scheme was proposed and analyzed in [11]. We describe this estimation scheme below. Recall that $SR_{\lambda,l}(X)$ is the solution to the following stochastic problem:

$$\text{minimize } t, \quad \text{subject to } \mathbb{E}[l(-X-t)] \leq \lambda. \tag{9}$$

Since we do not have access to the true distribution of X, instead of solving eq. (9) we use m i.i.d samples $\{Z_i\}_{i=1}^m$ (also denoted as a random vector $\mathbf{Z}$) from X to solve the following problem to estimate $SR_{\lambda,l}(X)$:

$$\text{minimize } t, \quad \text{subject to } \frac{1}{m}\sum_{i=1}^{m} l(-Z_i-t) \leq \lambda. \tag{10}$$

For $m \geq 1$, define the function $SR_m : \mathbb{R}^m \to \mathbb{R}$ as follows:

$$SR_m(z) \triangleq \min\left\{ t \in \mathbb{R} \,\middle|\, \frac{1}{m}\sum_{j=1}^{m} l(-z_j-t) \leq \lambda \right\}. \tag{11}$$

Then, $SR_m(\mathbf{Z})$ is the solution to eq. (10) and an estimator of $SR_{l,\lambda}(X)$. For the analysis of this UBSR estimator, we make the following assumption on the random variable X:

Assumption 6: $\exists q \geq 2, T > 0$ such that $\|X\|_{L^q}$ is finite and $\|X\|_{L^2} \leq T$.

The result below presents error bounds for the UBSR estimator $SR_m(\mathbf{Z})$.

Lemma 2: Let $\Delta_m \triangleq SR_{l,\lambda}(X) - SR_m(\mathbf{Z})$. Suppose Assumption 6 holds for $q > 4$ and assumptions 4 and 5 hold. Then, we have

$$\mathbb{E}\left[|\Delta_m|\right] \leq \frac{C_1}{\sqrt{m}}, \quad \text{and } \mathbb{E}\left[|\Delta_m|^2\right] \leq \frac{C_2}{\sqrt{m}},$$

where $C_1 = \frac{48L_1T}{b}$ and $C_2 = \frac{108L_1^2T^2}{b^2}$ respectively. See [12, Appendix II] for the proof.

The result below is a variant of Lemma 2 that caters to smooth loss functions.

Lemma 2′: Suppose Assumption 6 holds for $q > 8$ and assumption 4′ and assumption 5 hold. Then we have

$$\mathbb{E}\left[|\Delta_m|\right] \leq \frac{C_3}{m^{1/4}}, \text{ and } \mathbb{E}\left[|\Delta_m|^2\right] \leq \frac{C_4}{\sqrt{m}},$$

where $C_3 = \frac{(27L_2+8a)T^2}{b}$ and $C_4 = \frac{(270L_2^2+108a^2)T^4}{b^2}$. See [12, Appendix II] for the proof.

The constants $C_1, \ldots, C_4$ appearing in the lemmas above are derived based on a result in [16], and depend inversely on q, i.e., assuming a higher moment bound leads to lower constants. The conditions on q in both lemma 2 and lemma 2′ are required to obtain a tight bound on the Wasserstein distance between X and its empirical estimate Z. Interested readers are referred to [16] for more details.

While computing $SR_m(\mathbf{Z})$ requires solving a convex optimization problem, a closed form expression is not available for any given loss function. Instead, it is possible to obtain an estimator within δ-neighbourhood of $SR_m(\mathbf{Z})$ for any $\delta > 0$ using bisection search. For details on the bisection search algorithm, refer [12].

V. UBSR Optimization

Let $\Theta \subset \mathbb{R}^d$ be a compact and convex set. Given a function $F : \Theta \times \mathbb{R} \to \mathbb{R}$ and a random variable ξ, a standard stochastic optimization algorithm deals with the problem of minimizing $\mathbb{E}[F(\theta,\xi)]$ using samples of either $F(\theta,\xi)$ or $\nabla_\theta F(\theta,\xi)$. Instead, we are interested in the problem of minimizing the UBSR of $F(\theta,\xi)$, i.e., to find a

$$\theta^* \in \operatorname*{arg\,min}_{\theta\in\Theta} h(\theta), \text{ where } h(\theta) \triangleq SR_{l,\lambda}(F(\theta,\xi)). \tag{12}$$

In our setting, we can obtain samples of ξ, which can be used to compute $F(\theta,\xi)$ and $\nabla_\theta F(\theta,\xi)$. Under the assumption that for every $\theta \in \Theta, F(\theta,\xi) \in \mathcal{X}_l$, we express $h : \Theta \to \mathbb{R}$ as follows[1]:

$$h(\theta) = \inf\left\{ t \in \mathbb{R} \,|\, \mathbb{E}[l(-F(\theta,\xi)-t)] \leq \lambda \right\}.$$

In the next section, we derive properties of UBSR and establish conditions under which h is strongly convex.

A. Properties of UBSR

For deriving the UBSR gradient expression, we require Assumptions 1 to 3 to hold for each random variable in $\{F(\theta,\xi) : \theta \in \Theta\}$. For the remainder of the paper, whenever any of these assumption holds, it is implied to hold for every random variable in $\{F(\theta,\xi) : \theta \in \Theta\}$. We first define $g : \mathbb{R} \times \Theta \to \mathbb{R}$ as follows:

$$g(t,\theta) \triangleq \mathbb{E}[l(-F(\theta,\xi)-t)] - \lambda. \tag{13}$$

Proposition 3.1 implies that for every $\theta \in \Theta$, $g(\cdot,\theta)$ is continuous, strictly decreasing, and has a unique root that coincides with $h(\theta)$, i.e., for every $\theta \in \Theta, g(h(\theta),\theta) = 0$.

We require the result above to invoke the implicit function theorem in the derivation of the UBSR gradient expression in the following section.

The result below shows that h is strongly convex under conditions similar to those in Proposition 3.2, except that we impose a strong concavity requirement on F. These assumptions are satisfied in the context of a mean-variance portfolio optimization problem [12, Section VI-B].

Lemma 3: Let l be convex and let $F(\cdot,\xi)$ be μ-strongly concave w.p. 1. Let Assumptions 1 to 3 hold, then h is μ-strongly convex.

[1]For notational convenience, we suppress the dependency of l, λ on h.

B. UBSR gradient and its estimation

We require the following assumption for deriving the UBSR gradient expression. A similar assumption has been made earlier in [11], [6] in the context of bounded random variables.

Assumption 7: $F(\cdot,\xi)$ is continuously differentiable almost surely.

We now present the main result that provides an expression for the gradient of UBSR.

Theorem 5.1 (Gradient of UBSR): Suppose Assumptions 1 to 3 and 7 hold. Then the function h is continuously differentiable and the gradient of h can be expressed as follows:

$$\nabla h(\theta) = -\frac{\mathbb{E}\Big[l'(-F(\theta,\xi)-h(\theta))\nabla F(\theta,\xi)\Big]}{\mathbb{E}\big[l'(-F(\theta,\xi)-h(\theta))\big]}. \tag{14}$$

See [12, Appendix III] for the proof. Further, the function h is Lipschitz, and we denote the associated Lipschitz constant with L_3.

From eq. (14), it is apparent that an estimate of $h(\theta)$ is required to form an estimate for $\nabla h(\theta)$. For estimating $h(\theta)$, we employ the scheme presented in Section IV. Suppressing the dependency on l,λ, we define the functions $SR_\theta^m : \mathbb{R}^m \to \mathbb{R}$ and $J_\theta^m : \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}^d$ as follows:

$$SR_\theta^m(\mathbf{z}) \triangleq \min\left\{ t\in\mathbb{R} \,\middle|\, \frac{1}{m}\sum_{j=1}^{m} l(-F(\theta,z_j)-t) \le \lambda \right\}, \tag{15}$$

$$J_\theta^m(\mathbf{z},\hat{\mathbf{z}}) \triangleq \frac{\sum_{j=1}^{m}\Big[l'(-F(\theta,z_j)-SR_\theta^m(\hat{\mathbf{z}}))\nabla F(\theta,z_j)\Big]}{\sum_{j=1}^{m}\big[l'(-F(\theta,z_j)-SR_\theta^m(\hat{\mathbf{z}}))\big]}. \tag{16}$$

Given a $\theta\in\Theta$, we use $SR_\theta^m(\cdot), J_\theta^m(\cdot,\cdot)$ to estimate $h(\theta), \nabla h(\theta)$ respectively. If $z,\hat{z}$ are constructed using i.i.d. samples of ξ, then one can obtain tight bounds on the estimation error, as derived in the following section. Note that the double sampling from ξ is necessary to avoid cross terms and has been used previously in [6].

To derive error bounds for eq. (16), it is necessary to bound the distance between UBSR estimate SR_θ^m and the true UBSR $h(\theta)$. This bound can be obtained from Lemma 2 by replacing Assumption 6 with the Assumption 8 below. To that end, we restate Lemma 2 as Lemma 4 below for the sake of readability.

Assumption 8: There exists $q\ge 2$ and $T>0$ such that for every $\theta\in\Theta$, $\|F(\theta,\xi)\|_{L^q}$ is finite and $\|F(\theta,\xi)\|_{L^2}\le T$.

Lemma 4 (UBSR estimation bounds): Suppose Assumptions 4, 5 and 8 hold and $\mathbf{Z}$ be an m-dimensional random vector such that each Z_j is an independent copy of ξ. Then,

$$\mathbb{E}[|SR_\theta^m(\mathbf{Z})-h(\theta)|] \le \frac{C_1}{\sqrt{m}},\ \mathbb{E}\left[[SR_\theta^m(\mathbf{Z})-h(\theta)]^2\right] \le \frac{C_2}{m}, \tag{17}$$

where C_1, C_2 are given in Lemma 2.

We make the following assumptions on F and ∇F.

Assumption 9: Suppose $p\in[1,\infty]$ is such that the p^{th} moments of $F, \nabla F$ exist and are finite, and there exists $M>0, L_4>0, L_5>0$ such that for every $\theta\in\Theta$, we have $\|\nabla F(\theta,\xi)\|_{L^p}\le M$, and for every $\theta_1,\theta_2\in\Theta$,

$$\|F(\theta_1,\xi)-F(\theta_2,\xi)\|_{L^p} \le L_4\|\theta_1-\theta_2\|_p,$$
$$\|\nabla F(\theta_1,\xi)-\nabla F(\theta_2,\xi)\|_{L^p} \le L_5\|\theta_1-\theta_2\|_p.$$

Similar assumptions have been made before in [6] for the non-asymptotic analysis of UBSR optimization scheme. Under these assumptions we establish that the objective function $h(\cdot)$ is smooth in the next section.

We now derive error bounds on the gradient estimator J_θ^m, under the following bounded variance assumption:

Assumption 10: There exist $\sigma_1,\sigma_2>0$ such that the following bounds hold for every $\theta\in\Theta$ and $i\in\{1,2,\dots,d\}$:

$$\mathbf{Var}(l'(-F(\theta,\xi)-h(\theta))\frac{\partial F(\theta,\xi)}{\partial\theta_i}) \le \sigma_1^2,$$
$$\mathbf{Var}(l'(-F(\theta,\xi)-h(\theta))) \le \sigma_2^2.$$

An assumption that bounds the variance of the gradient estimate is common to the non-asymptotic analysis of stochastic gradient algorithms, cf. [17], [18]. The assumption above is made in a similar spirit, and the result below establishes that the mean-squared error of the UBSR gradient estimate eq. (16) vanishes asymptotically at a $O(1/m)$ rate.

Lemma 5 (UBSR gradient estimator bounds): Suppose Assumptions 1 to 5 and Assumptions 7 to 10 hold. Let $\mathbf{Z},\hat{\mathbf{Z}}$ denote m-dimensional random vectors such that each Z_j and each $\hat{Z}_j$ are i.i.d copies of ξ, and, $\mathbf{Z}$ and $\hat{\mathbf{Z}}$ are independent. Then for every $\theta\in\Theta$, the gradient estimator $J_\theta^m(\mathbf{Z},\hat{\mathbf{Z}})$ defined in eq. (16) satisfies

$$\mathbb{E}\left[\left\|J_\theta^m(\mathbf{Z},\hat{\mathbf{Z}})-\nabla h(\theta)\right\|_1\right] \le \frac{D_1}{\sqrt{m}},$$
$$\text{and}\quad \mathbb{E}\left[\left\|J_\theta^m(\mathbf{Z},\hat{\mathbf{Z}})-\nabla h(\theta)\right\|_2^2\right] \le \frac{D_2}{m},$$

where $D_1 = \frac{2C_1L_1L_2M+L_1(\sigma_1\sqrt{d}+M\sigma_2)}{b^2}$ and $D_2 = \frac{8C_2L_1^2L_2^2M^2+4L_1^2(d\sigma_1^2+M^2\sigma_2^2)}{b^4}$.

See [12, Appendix III] for the proof. The above result shows that to get ϵ-accurate gradient estimate in the RMSE sense, one requires samples of the order of $\mathcal{O}(1/\epsilon^2)$. A similar proof leads to the MSE bound of the order $\mathcal{O}(1/\sqrt{m})$ for non-Lipschitz loss functions, that involves invoking lemma 2' instead of lemma 2.

C. SG Algorithm for UBSR optimization

In each iteration k of this algorithm, we sample m_k-dimensional random vectors $\mathbf{Z^k}, \hat{\mathbf{Z}}^\mathbf{k}$ that are independent of one another and independent of the previous samples, such that for every $i\in[1,2,\dots,m_k]$, $Z_i^k\sim\xi, \hat{Z}_i^k\sim\xi$. For the SG algorithm to optimize UBSR, one can perform the following update with an arbitrarily chosen $\theta_0\in\Theta$:

$$\theta_k = \Pi_\Theta\left(\theta_{k-1}-\alpha_k J^k\right), \tag{18}$$

where $\Pi_\Theta(x)\triangleq \arg\min_{\theta\in\Theta}\|x-\theta\|_2$ denotes the operator that projects onto the convex and compact set Θ.

Let $\mathcal{F}_0 = \sigma(\theta_0)$ and for every $k \in \mathbb{N}$, let $\mathcal{F}_k = \sigma\left(\theta_0, \mathbf{Z^1}, \hat{\mathbf{Z}}^\mathbf{1}, \ldots, \mathbf{Z^k}, \hat{\mathbf{Z}}^\mathbf{k}\right)$. Then $\{\mathcal{F}_k\}_{k\geq 0}$ forms the filtration that θ_n is adapted to. Applying the independence lemma from [19] with Lemma 5, we have $\forall k \in \mathbb{N}$,

$$\mathbb{E}\left[\left\|\hat{J}_\theta^{m_k}(\mathbf{Z}) - \nabla h(\theta_{k-1})\right\|_1 \middle| \mathcal{F}_{k-1}\right] \leq \frac{\hat{D}_1}{\sqrt{m_k}},$$
$$\mathbb{E}\left[\left\|\hat{J}_\theta^{m_k}(\mathbf{Z}) - \nabla h(\theta_{k-1})\right\|_2^2 \middle| \mathcal{F}_{k-1}\right] \leq \frac{\hat{D}_2}{m_k}.$$

For the algorithm presented above, we derive non-asymptotic bounds for two choices of the batch-sizes, namely constant and increasing.

Theorem 5.2: Let Assumptions 1 to 5 and Assumptions 7 to 10 hold, and let h be μ-strongly convex. Let θ^* be the minimizer of $h(\cdot)$ and let $z_n \triangleq \theta_n - \theta^*$. Let $\alpha_k = \frac{c}{k}$, with c satisfying $1 \leq \mu c \leq 3$. If $m_k = k$, $\forall k$, then

$$\mathbb{E}\left\|z_n\right\|_2^2 \leq \frac{512c^2\hat{D}_1^2}{n+1} + \frac{450\mathbb{E}\left\|z_0\right\|_2^2 + 128c^2\hat{D}_2\ln(n)}{(n+1)^2},$$

where $\hat{D}_1 = \frac{2(C_1+d_1)L_1L_2M + L_1(\sigma_1\sqrt{d}+M\sigma_2)}{b^2}$ and $\hat{D}_2 = \frac{8(C_2+d_1^2)L_1^2L_2^2M^2 + 4L_1^2(d\sigma_1^2+M^2\sigma_2^2)}{b^4}$.

In addition, if $m_k = m$ for all k, then we have

$$\mathbb{E}\left[\left\|z_n\right\|_2^2\right] \leq \frac{64c^2\hat{D}_1^2}{m} + \frac{450\mathbb{E}\left[\left\|z_0\right\|_2^2\right] + 64c^2\hat{D}_2}{(n+1)m}.$$

See [12, Appendix III] for the proof.

Corollary 5.1: Under conditions of Theorem 5.2, h is L_6-smooth, and

$$\mathbb{E}\left[h(\theta_n) - h(\theta^*)\right] \leq L_6\mathbb{E}\left[\left\|\theta_n - \theta^*\right\|_2^2\right],$$

where $L_6 = \frac{L_2M\left(\frac{1+L_1(L_3+L_4)}{b}\right) + L_1L_5}{b}$.

Remark 2: Asymptotic convergence rate of $\mathcal{O}(1/n)$ has been derived earlier in [11] for the scalar UBSR optimization case, but their result required a batchsize $m \geq n$ for each iteration. Our result not only establishes a non-asymptotic bound of the same order, but also allows for an increasing batchsize that does not depend on n. Table I summarizes the convergence rates for different choices of batch-sizes.

Remark 3: For the case of a non-Lipschitz loss function, we can obtain a non-asymptotic bound of $\mathcal{O}(1/\sqrt{n})$ by replacing assumption 4 with assumption $4'$. Owing to space constraints, we omit the details.

Remark 4: The bound in Theorem 5.2 can be inferred for the special case where the loss function is $l(x) = \exp(\beta x)$, $\lambda = 1$, and the underlying distribution is Gaussian. In this case, UBSR reduces to the entropic risk measure, and the loss function l is neither strongly convex nor smooth. However, using a closed-form expression of UBSR (see the portfolio optimization example in Section VI-B of [12]), the strong convexity and smoothness of h can be established directly. Such strong convexity and smoothness results can be used to arrive at a bound similar to that in Theorem 5.2 by employing completely parallel arguments in the proof.

TABLE I

COMPLEXITY BOUNDS FOR UBSR-SG TO ENSURE $\mathbb{E}\left[h(\theta_n) - h(\theta^*)\right] \leq \epsilon$.

Batchsize	$m_k = k$	$m_k = n^p$
Iteration complexity	$\mathcal{O}(1/\epsilon)$	$\mathcal{O}\left(1/\epsilon^{\frac{1}{p}}\right)$
Sample complexity	$\mathcal{O}(1/\epsilon^2)$	$\mathcal{O}\left(1/\epsilon^{1+\frac{1}{p}}\right)$

VI. CONCLUSIONS

We lay the foundations for UBSR estimation and optimization for the case of unbounded random variables. Our contributions are appealing in financial applications such as portfolio optimization as well as risk-sensitive reinforcement learning. As future work, it would be interesting to (i) explore UBSR optimization in the non-convex case; and ii) develop Newton-based methods for UBSR optimization.

REFERENCES

[1] P. Jorion, *Value at Risk: The New Benchmark for Controlling Market Risk*. McGraw-Hill, 1997.

[2] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42, 2000.

[3] C. Acerbi and D. Tasche, "On the coherence of expected shortfall," *Journal of Banking & Finance*, vol. 26, no. 7, pp. 1487–1503, 2002.

[4] H. Föllmer and A. Schied, *Stochastic finance*. de Gruyter, 2016.

[5] J. Dunkel and S. Weber, "Stochastic root finding and efficient estimation of convex risk measures," *Operations Research*, vol. 58, no. 5, pp. 1505–1521, 2010.

[6] V. Hegde, A. S. Menon, L. A. Prashanth, and K. Jagannathan, "Online Estimation and Optimization of Utility-Based Shortfall Risk," *Mathematics of Operations Research*, 2024 (to appear).

[7] L. A. Prashanth and S. P. Bhat, "A Wasserstein Distance Approach for Concentration of Empirical Risk Estimates," *Journal of Machine Learning Research*, vol. 23, no. 238, pp. 1–61, 2022.

[8] H. Föllmer and A. Schied, *Stochastic Finance: An Introduction in Discrete Time*. Berlin, Boston: De Gruyter, 2004.

[9] K. Giesecke, T. Schmidt, and S. Weber, "Measuring the risk of large losses," *Journal of Investment Management*, vol. 6, 01 2008.

[10] S. Weber, "Distribution-invariant risk measures, information, and dynamic consistency," *Mathematical Finance*, vol. 16, no. 2, pp. 419–441, 2006.

[11] Z. Hu and D. Zhang, "Utility-based shortfall risk: Efficient computations via monte carlo," *Naval Research Logistics (NRL)*, vol. 65, no. 5, pp. 378–392, 2018.

[12] S. Gupte, L. A. Prashanth, and S. P. Bhat, "Optimization of utility-based shortfall risk: A non-asymptotic viewpoint," *arXiv preprint arXiv:2310.18743*, 2023.

[13] P. Artzner, F. Delbaen, J. Eber, and D. Heath, "Coherent measures of risk," *Mathematical Finance*, vol. 9, no. 3, pp. 203–228, 1999.

[14] A. Ahmadi-Javid, "Entropic value-at-risk: A new coherent risk measure," *Journal of Optimization Theory and Applications*, vol. 155, no. 3, pp. 1105–1123, 2012.

[15] A. J. Kleywegt, A. Shapiro, and T. H. de Mello, "The sample average approximation method for stochastic discrete optimization," *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2002.

[16] Fournier, Nicolas, "Convergence of the empirical measure in expected wasserstein distance: non-asymptotic explicit bounds in rd," *ESAIM: PS*, vol. 27, pp. 749–775, 2023.

[17] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *Siam Review*, vol. 60, no. 2, pp. 223–311, 2018.

[18] N. Bhavsar and L. A. Prashanth, "Nonasymptotic bounds for stochastic optimization with biased noisy gradient oracles," *IEEE Transactions on Automatic Control*, vol. 68, no. 3, pp. 1628–1641, 2023.

[19] S. E. Shreve, *Stochastic Calculus for Finance I: The Binomial Asset Pricing Model*, ser. Springer Finance. Springer New York, NY, 2004.