Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

DATA-DRIVEN SOLUTIONS AND UNCERTAINTY QUANTIFICATION FOR MULTISTAGE STOCHASTIC OPTIMIZATION

Yunhao Yan[1], and Henry Lam[2]

[1,2]Dept. of Industrial Eng. and Operations Research, Columbia University, New York, NY, USA

ABSTRACT

Multistage stochastic optimization problems appear commonly in various disciplines in operations management and system control. While computational methods have been studied, they are mostly restricted to two-stage problems or otherwise require strong assumptions. In this paper, we propose a simple data-driven solution technique and uncertainty quantification method for these problems, based on natural multisample generalizations of the well-known sample average approximation and the so-called single replication procedure. Under the assumptions of stagewise independence and the use of parameterized policies, we justify statistical consistency and a coverage guarantee on bounding the optimality gap using our approaches. Our developments entail the establishment of several new statistical properties of the so-called multisample U-process that closely connect to multistage stochastic optimization.

1 INTRODUCTION

A multistage stochastic optimization problem comprises a sequential decision process of an agent over a given time horizon, where each time or stage entails a new randomness that potentially depends on the past. Correspondingly, the agent makes a decision at each stage depending on past history, the collection of which is often known as a policy. The goal is to devise a policy that minimizes the overall expected penalty (or equivalently maximizes the expected reward). This problem appears commonly in various disciplines such as operations management (Talluri and Van Ryzin 2006; Scarf 1960) and stochastic control (Koussoulas and Leondes 1986; Moore, Zhou, and Lim 1999).

In this paper, we are interested in data-driven solution techniques and uncertainty quantification for multistage stochastic optimization. While computational methods have been studied for such problems in the past (Swamy and Shmoys 2005; Shapiro 2003), they appear mostly restricted to two-stage problems or otherwise require strong assumptions. Part of the challenges seem to be the complexity involved in extending single- or two-stage techniques to more general sequential settings (Shapiro and Nemirovski 2005; Shapiro 2008). Because of this, the asymptotics of SAA (Shapiro, Dentcheva, and Ruszczynski 2021 Chapter 5), and inference methods such as the batching (Mak, Morton, and Wood 1999), the so-called single replication procedure (SRP) (Bayraksan and Morton 2006), the bootstrap (Eichhorn and Römisch 2007; Chen and Woodruff 2023) and bagging (Lam and Qian 2018a; Chen and Woodruff 2023; Lam and Qian 2018b) are all studied in the contexts of single- or two-stage problems. For general multi-stage problems, the complexity of SAA scenarios increases exponentially with the time horizon, thus rendering these techniques more challenging to use and provide guarantees.

Motivated by the above, our goal in this paper is to make a step towards the solution and inference of multistage stochastic optimization problems. In particular, we consider two assumptions. First is *stagewise independence*, meaning that the randomness at each stage is independent of the past. This assumption has

been used in Shapiro, Dentcheva, and Ruszczynski (2021) Section 5.8 which provides some consistency results for multistage SAA, but under several additional strong assumptions. Second is *parameterized policies*, meaning that instead of searching for a globally optimal policy we restrict our policy class to a parametric form, thereby reducing the potentially infinite-dimensional problem into optimization over the Euclidean space. These two assumptions are motivated from the gain on tractability in handling multistage optimization. While they result in a loss on generality, there are important problems that fall under these assumptions, as we will describe further in the sequel.

More precisely, under these assumptions, we propose a natural multisample generalization of the well-known sample average approximation (SAA) method to integrate data into multistage stochastic optimization. We call our approach *multisample average approximation (MSAA)*. Correspondingly, we also devise a method to construct statistical confidence bounds for the optimality gap of a given solution, based on a natural generalization of SRP. This procedure, proposed in Bayraksan and Morton (2006) for single or at most two-stage problems, entails a formula that resembles the normality confidence bound, but applies even to problems that exhibit non-normal limits (such as problems with multiple optimal solutions) and requires an insightful analysis that deviates from the standard route. We will present a multisample adaptation of this procedure, which we call multisample single replication procedure (MSRP), and demonstrate how a similar analysis as SRP can apply to multistage problems.

We present statistical consistency of our MSAA and the asymptotic coverage guarantee of MSRP. Technically, with the stagewise independence and parameterized policy assumptions, the multistage stochastic optimization problem reduces into the optimization of a so-called *multisample U-process*, which is a statistical notion that we leverage heavily and makes our analysis tractable. U-statistic can be viewed as a generalization of the sample mean but, instead of having one observation at each summand in the averaging, it uses a symmetric function of several different observations in the summand, and multisample U-statistic further generalizes to several independent data sources. The theory of multisample U-statistics dates back to seminal works such as Lehmann (1951) and Dwass (1956), with convergence properties examined in Sen (1974) and Gut (1976), and further exploration in de la Pena (1992) and Arcones and Giné (1993). We harness some of these tools and, along our analyses, we develop several new theoretical results concerning the statistical properties of multsample U-statistics.

The rest of the paper is organized as follows. In Section 2, we introduce multistage stochastic optimization, describe our assumptions on stagewise independence and parameterized policy, and provide a few examples. In Section 3, we review several asymptotic results on multisample U-statistics that we will utilize to analyze our solution and uncertainty quantification methods. In Section 4, we present statistical consistency results on MSAA. In Section 5, we devise our MSRP to bound optimality gaps and derive its asymptotic coverage validity.

2 MULTISTAGE STOCHASTIC OPTIMIZATION

Consider a multistage stochastic programs described as follows. Denote a sequence of independent random variables $\xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(c)}$ where each of them is realized sequentially at each time over a given time horizon c. At each time $k = 1, 2, \ldots, c$, given $\xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(k-1)}$ (with no observation at $k = 1$), the agent makes a decision $a_k \in D_k$, where D_k is some decision space, and observe the realized random variable $\xi^{(k)}$. Let $h(a_1, a_2, \ldots, a_{c-1}; \xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(c)})$ be a penalty function. Now denote the sample space of the first k variables, i.e. $(\xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(k)})$, to be $\mathscr{S}_k$. A decision rule at time k is defined as a deterministic map $A_k : \mathscr{S}_{k-1} \to D_k$, i.e. $a_k = A_k(\xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(k-1)})$. We denote the space of all feasible decision rules at time k to be $\mathscr{A}_k$., Correspondingly, a policy is defined as the combination of $c-1$ decision rules

across $k = 1,2,\ldots,c-1$, i.e. $(A_1,\ldots,A_{c-1})$, and the policy space is $\prod_{k=1}^{c-1}\mathscr{A}_k$. With these, we consider the multistage stochastic program

$$\begin{aligned} \textbf{minimize} \quad & \mathbb{E}[h(A_1, A_2(\xi^{(1)}),\ldots,A_{c-1}(\xi^{(1)},\ldots,\xi^{(c-1)});\xi^{(1)},\xi^{(2)},\ldots,\xi^{(c)})] \\ \textbf{subject to} \quad & (A_1,\ldots,A_{c-1}) \in \prod_{k=1}^{c-1}\mathscr{A}_k. \end{aligned} \tag{1}$$

We consider a parametrized policy class. That is, $A_k \in \mathscr{A}_k$ takes the form $A_k(\xi^{(1)},\xi^{(2)},\ldots,\xi^{(k)}) = g(\theta_k;\xi^{(1)},\xi^{(2)},\ldots,\xi^{(k)})$ where θ_k is a lower-dimensional parameter than A_k itself. Typically, θ_k for $k = 1,2,\ldots,c$ are in the Euclidean space. With this, the stochastic program (1) is reduced from a potentially infinite-dimensional space to a finite dimension. Now, replacing the decision variables a_k with the deterministic functions $g(\theta_k;\xi^{(1)},\xi^{(2)},\ldots,\xi^{(k)})$, we can simplify (1) as

$$\underset{\theta\in\Theta}{\textbf{minimize}} \quad \mu(\theta) \coloneqq \mathbb{E}[h(\theta;\xi^{(1)},\xi^{(2)},\ldots,\xi^{(c)})] \tag{2}$$

where $\Theta \subset \mathbb{R}^d$ denotes a subset in the Euclidean space of dimension d.

Furthermore, note that we have assumed above that $\xi^{(1)},\xi^{(2)},\ldots,\xi^{(c)}$ are independent, or in other words stagewise independence. Suppose each $\xi^{(k)}$ is drawn from distribution F_k, for $k = 1,2,\ldots,c$. Now, consider the settings where there are N_k i.i.d. observations for each distribution F_k, denoted as $\{\xi_{i_k}^{(k)}\}_{i_k=1}^{N_k}$. Based on these samples, we can formulate an empirical counterpart of (2), by plugging in the empirical distribution for each time k, i.e.,

$$\underset{\theta\in\Theta}{\textbf{minimize}} \qquad \mu_{\mathbf{N}}(\theta) \coloneqq \left(\prod_{k=1}^{c}\frac{1}{N_k}\right)\sum_{i_1=1}^{N_1}\sum_{i_2=1}^{N_2}\cdots\sum_{i_c=1}^{N_c} h(\theta;\xi_{i_1}^{(1)},\xi_{i_2}^{(2)},\ldots,\xi_{i_c}^{(c)}) \tag{3}$$

where the subscript $\mathbf{N} \coloneqq (N_1,N_2,\ldots,N_k)$ represents the vector of sample sizes. We refer to this formulation as Multisample Average Approximation (MSAA). When $\mathbf{N}$ is sufficiently large for each component, solving the MSAA problem is expected to give rise to a nearly optimal solution for the original problem (2).

Our interest in the MSAA is primarily based on two considerations. Firstly, the true distribution F_k may not be known, but we may have real-world data (such as sales figures or customer records in the examples below). Secondly, even with perfect information of all distributions F_k, solving the stochastic program 2 presents challenges when dealing with continuous distributions, and the finite-sample approximation through MSAA may allow for efficient computation. Note that our MSAA relies critically on the assumptions of stagewise independence and parameterized policy. While they clearly lead to a loss in generality, there are still important problems that fall under these conditions. We introduce several examples.

Example 1 (Airline Management) We consider the c-class airline capacity control model in Talluri and Van Ryzin (2006) Chapter 2.2.2. In this model, an agent aims to maximize revenue from selling a total of C airline tickets. For $k = 1,2,\ldots,c$, $\xi^{(k)}$ customers from the k-th class arrive in order and express willingness to purchase tickets at price p_k. It is a conventional assumption that $p_1 < p_2 < \ldots < p_c$, as higher-value customers, being more time-sensitive, tend to book flights closer to the departure date. The decision variable for the agent is $\theta \coloneqq (\theta_1,\theta_2,\ldots,\theta_{c-1}) \in [0,C]^{c-1}$, where θ_k represents the number of tickets reserved at time k for later-arriving, higher-valued customers. Consequently, the total revenue is a deterministic function of both θ and the realized arrivals $\{\xi^{(k)}\}_{k=1}^{c}$. Thus the problem can be formulated as a stochastic program of the form (2).

In the above setup, other than airline ticket sales which must be integers, all variables including capacities and decisions can be treated as continuous. Talluri and Van Ryzin (2006) present two methods for solving the capacity control problem: dynamic programming with discretization of θ and Monte Carlo integration. The latter method is essentially the SAA approach. The Monte Carlo approach is often preferred due to its computational efficiency. The time complexity for dynamic programming is $O\left(cC^2/\varepsilon^2\right)$, where ε denotes the discretization precision. However, an analysis on the optimality guarantee for the Monte Carlo method is lacking in the literature, an issue we address in this paper.

Example 2 (Inventory Management with (s,S)-policy) A single item inventory management problem involves managing the amount I_k of a resource using a replenishment policy, where the main tradeoff lies between storage and backorder costs. Let I_k represent the inventory level at beginning of time k. The agent first decides on the replenishment amount R_k (which could be zero) at a cost $C(R_k)$. Subsequently, the demand D_k is reviewed. Here the stagewise independence of D_k is assumed. If $I_k+R_k-D_k>0$, a holding cost h_k per unit is incurred; if $I_k+R_k-D_k<0$, a backorder cost b_k per unit is incurred. The total management cost over the horizon $k=1,2,\ldots,c$ is given by $\sum_{k=1}^{c}(I_k+R_k-D_k)^+h_k+(D_k-I_k-R_k)^+b_k+C(R_k)$. Although the replenishment policy can be any arbitrary function of the tuple (k,I_k), in many cases, such as when the replenishment cost is linear, i.e., $C(R_k)=C_0+C_kR_k$ if $R_k>0$ and no cost if $R_k=0$, the optimal policy is an (s_k,S_k)-policy as shown by Scarf (1960). Under this policy, if $I_k<s_k$, the inventory is replenished up to $S_k>s_k$. This insight simplifies the optimization problem from an infinite-dimensional space of decision functions to a Euclidean space, represented by $\theta:=(s_k,S_k)_{k=1}^{c}$ in $\mathbb{R}^{2c}$. Consequently, this reduction allows the problem to be formulated as in (2), making it suitable for the application of MSAA.

Example 3 (Discrete Linear-Quadratic Stochastic Control) As studied by Koussoulas and Leondes (1986) and Moore, Zhou, and Lim (1999), Linear-Quadratic Control (LQC) aims to minimize the quadratic norm of a vector $\mathbf{x}k\in\mathbb{R}^s$ subject to linear dynamics with random noise over a time horizon c. A distinctive aspect of LQC is that the optimal control variable $u_k\in\mathbb{R}^d$ employs a linear feedback control policy from the state variable x_k, specifically, $u_k=M_kx_k$ where $M_k\in\mathbb{R}^{d\times s}$. With boundedness assumptions, the parameters $M_kk=1^c$ are confined within a closed set. Consequently, the problem reduces to the stochastic program in the form of (2) where MSAA can be applied.

3 BACKGROUND ON MULTISAMPLE U-STATISTICS

For a specified θ, the objective function $\mu_{\mathbf{N}}(\theta)$ in (3) constitutes a so-called *multisample U-statistic*, also referred to as a *generalized U-statistic* in some literature. A U-statistic is a generalization of sample mean where instead of having one observation in each summand, it consists of a symmetric function of multiple observations in each summand. A multisample U-statistic further generalizes U-statistic to allow for several independent sources of data.

With the above, MSAA can be viewed as a minimization of a multisample U-process, with data size $\mathbf{N}$, where the process refers to its dependence on θ. Given that the multisample U-statistic is less well-known than the U-statistic, we will introduce related definitions in Section 3.1, and discuss some classical asymptotic properties in Section 3.2 that are useful for our subsequent developments.

3.1 Multisample U-statistics

In order to introduce a multisample U-statistic, we first define the so-called kernel function. Let $\mathbf{m}=(m_1,m_2,\ldots,m_c)$ denote a c-dimensional vector of positive integers and $|\mathbf{m}|=\sum_{k=1}^{c}m_k$. Also let $\mathbf{x}^{(k)}=(x_1^{(k)},\ldots,x_{m_k}^{(k)})\in\mathbb{R}^{m_k}$.

Definition 1 (Kernel Function) $h_{\mathbf{m}}(\mathbf{x}^{(1)};\mathbf{x}^{(2)};\ldots;\mathbf{x}^{(c)}):\mathbb{R}^{|\mathbf{m}|}\to\mathbb{R}$ is called a kernel function with index $\mathbf{m}$ if it is symmetric to the interchange of variables within each segment k.

For example, when $\mathbf{m}=(2,2)$, $h_{\mathbf{m}}(x_1^{(1)},x_2^{(2)};y_1^{(1)},y_2^{(2)})=h_{\mathbf{m}}(x_2^{(1)},x_1^{(2)};y_1^{(1)},y_2^{(2)})=h_{\mathbf{m}}(x_1^{(1)},x_2^{(2)};y_2^{(1)},y_1^{(2)})$. However, it is prohibited to exchange arguments across segments of different $\mathbf{x}^{(k)}$, i.e. $h_{\mathbf{m}}(x_1^{(1)},x_2^{(2)};y_1^{(1)},y_2^{(2)})\neq h_{\mathbf{m}}(y_1^{(1)},x_2^{(2)};x_1^{(1)},y_2^{(2)})$ in general.

For $k=1,2,\ldots,c$, let F_k be c distinct unknown distributions on $\mathbb{R}$. Now assume that we have N_k observations from F_k that are denoted as $\{\xi_{i_k}^{(k)}\}_{i_k=1}^{N_k}$. Recall that $\mathbf{N}=(N_1,N_2,\ldots,N_c)$ denotes the sample size vector. We have the following definition:

Definition 2 (Multisample U- and V-Statistics) Given $\mathbf{m}$ and $\{F_k\}_{k=1}^c$ as well as the i.i.d. random variables $\{\xi_{i_k}^{(k)}\}_{i_k=1}^{N_k}$ drawn from the distribution F_k, the multisample V-statistic and the multisample U-statistic with kernel function $h_{\mathbf{m}}$ and sample size $\mathbf{N}=(N_1,N_2,\ldots,N_c)$ are defined as

$$V_{\mathbf{N}}[h_{\mathbf{m}}]:=\left(\prod_{k=1}^{c}\frac{1}{N_k^{m_k}}\right)\left(\sum_{i_1^{(1)}=1}^{N_1}\cdots\sum_{i_{m_1}^{(1)}=1}^{N_1}\right)\cdots\left(\sum_{i_1^{(c)}=1}^{N_c}\cdots\sum_{i_{m_c}^{(c)}=1}^{N_c}\right)h_{\mathbf{m}}(\xi_{i_1^{(1)}}^{(1)},\ldots,\xi_{i_{m_1}^{(1)}}^{(1)};\ldots;\xi_{i_1^{(c)}}^{(c)},\ldots,\xi_{i_{m_c}^{(c)}}^{(c)}),\quad\text{and}$$

$$U_{\mathbf{N}}[h_{\mathbf{m}}]:=\left(\prod_{k=1}^{c}\frac{(N_k-m_k)!}{N_k!}\right)\left(\sum_{i_1^{(1)}=1}^{N_1}\cdots\sum_{i_{m_1}^{(1)}=1}^{N_1}\right)\cdots\left(\sum_{i_1^{(c)}=1}^{N_c}\cdots\sum_{i_{m_c}^{(c)}=1}^{N_c}\right)h_{\mathbf{m}}(\xi_{i_1^{(1)}}^{(1)},\ldots,\xi_{i_{m_1}^{(1)}}^{(1)};\ldots;\xi_{i_1^{(c)}}^{(c)},\ldots,\xi_{i_{m_c}^{(c)}}^{(c)})\times\mathbb{1}_{\text{dist.}},$$

where the indicator $\mathbb{1}_{\text{dist.}}:=\mathbb{1}\{\text{For each } k=1,\ldots,c,\ i_s^{(k)} \text{ are distinct for all } s=1,\ldots,m_k\}$ excludes the summation over diagonal terms.

Note that U- and V-statistics are similar except that U does not contain repeated observations in each summand. To analyze their asymptotics, we first notice that it is natural to understand the multisample U-statistic as an estimator for the expectation $\mathbb{E}[h_{\mathbf{m}}(\xi^{(1)};\xi^{(2)};\ldots;\xi^{(k)})]$, where $\xi^{(k)}=(\xi_1^{(k)},\xi_2^{(k)},\ldots,\xi_{m_k}^{(k)})$ are m_k i.i.d. from F_k and all $\{\xi^{(k)}\}_{k=1}^c$ are independent. To analyze convergence properly, note that the U-statistic $U_{\mathbf{N}}[h_{\mathbf{m}}]$ and the V-statistic $V_{\mathbf{N}}[h_{\mathbf{m}}]$ are defined over a c-dimensional array indexed by $\mathbf{N}=(N_1,N_2,\ldots,N_c)$. To proceed, we have to define the convergence of an array.

Definition 3 (Stochastic Convergence of Arrays) Let $S_{\mathbf{N}}$ be an c-dimensional array of random variables indexed by integer vector $\mathbf{N}=(N_1,N_2,\ldots,N_c)\geq 0$ and S be a given variable. The partial order $\mathbf{N}\geq\mathbf{N}'$ is defined in the elementwise sense. We define three modes of convergence as follows: (i) $S_{\mathbf{N}}\xrightarrow{\mathbb{P}}S$ if for any $\varepsilon,\delta>0$, there exists an $\mathbf{N}'$ such that $\sup_{\mathbf{N}\geq\mathbf{N}'}\mathbb{P}(|S_{\mathbf{N}}-S|>\varepsilon)<\delta$; (ii) $S_{\mathbf{N}}\xrightarrow{L^p}S$ if for any $\varepsilon>0$, there exists an $\mathbf{N}'$ such that $\sup_{\mathbf{N}\geq\mathbf{N}'}\mathbb{E}|S_{\mathbf{N}}-S|^p<\varepsilon$; (iii) $S_{\mathbf{N}}\xrightarrow{a.s.}S$ if for a.s. $\omega\in\Omega$ and $\varepsilon>0$, there exists an $\mathbf{N}'$ such that $\sup_{\mathbf{N}\geq\mathbf{N}'}|S_{\mathbf{N}}(\omega)-S(\omega)|<\varepsilon$.

In fact, either $S_{\mathbf{N}}\xrightarrow{L^p}S$ or $S_{\mathbf{N}}\xrightarrow{a.s.}S$ implies $S_{\mathbf{N}}\xrightarrow{\mathbb{P}}S$, the same as when $c=1$. The a.s.-convergence of a stochastic array is more complicated than that of a single sequence. It is usually challenging, if not possible, to achieve a.s.-convergence with mere integrability. In the literature, an alternative sufficient condition for a c-dimensional array is $E|X|(\log^+|X|)^{c-1}<\infty$ (Smythe 1973; Gut 1976; Sen 1977). When $c=2$, the specified condition corresponds to the Orlicz space $L\log^+L$. Therefore, we shall use the term *Orlicz condition* to refer to such conditions (for all $c\geq 2$). Recall the variable $h_{\mathbf{m}}(\xi^{(1)};\xi^{(2)};\ldots;\xi^{(k)})$, where $\xi^{(k)}=(\xi_1^{(k)},\xi_2^{(k)},\ldots,\xi_{m_k}^{(k)})$ are i.i.d. from F_k and all $\{\xi^{(k)}\}_{k=1}^c$ are independent. The subsequent

kernel types are relevant when discussing the convergence of $U_{\mathbf{N}}[h_{\mathbf{m}}]$ and $V_{\mathbf{N}}[h_{\mathbf{m}}]$ to the mean value $\mathbb{E}[h_{\mathbf{m}}(\xi^{(1)};\xi^{(2)};\ldots;\xi^{(k)})]$.

Definition 4 (Kernel Types) Let $\mathbf{F}_c \coloneqq F_1 \times F_2 \times \ldots \times F_c$ and $h_{\mathbf{m}}$ be given. Then

(i) $h_{\mathbf{m}} \in L^p(\mathbf{F}_c)$ if $\mathbb{E}|h_{\mathbf{m}}(\xi^{(1)};\xi^{(2)};\ldots;\xi^{(c)})|^p < \infty$;

(ii) $h_{\mathbf{m}} \in O^p(\mathbf{F}_c)$ if $\mathbb{E}[|h_{\mathbf{m}}(\xi^{(1)};\ldots;\xi^{(c)})|^p(\log^+|h_{\mathbf{m}}(\xi^{(1)};\ldots;\xi^{(c)})|^p)^{c-1}] < \infty$.

(iii) $h_{\mathbf{m}} \in \mathscr{L}^p(\mathbf{F}_c)$ or $h_{\mathbf{m}} \in \mathscr{O}^p(\mathbf{F}_c)$ if the respective condition for $L^p(\mathbf{F}_c)$ or $O^p(\mathbf{F}_c)$ holds for all $(\xi^{(1)};\ldots;\xi^{(c)})$ when the indices $i_1, i_2, \ldots, i_{m_k}$ in $\xi^{(k)} = (\xi^{(k)}_{i_1}, \ldots, \xi^{(k)}_{i_{m_1}})$ are not necessarily distinct.

To better clarify (iii), take for instance $h_{(\mathbf{2},\mathbf{2})}(x_1,x_2;y_1,y_2) \in \mathscr{L}^1(F_1 \times F_2)$. Then (iii) implies

$$\mathbb{E}|h(X_1,X_2;Y_1,Y_2)| + \mathbb{E}|h(X_1,X_1;Y_1,Y_2)| + \mathbb{E}|h(X_1,X_2;Y_1,Y_1)| + \mathbb{E}|h(X_1,X_1;Y_1,Y_1)| < \infty,$$

where X_i $(i = 1,2)$ and Y_i $(j = 1,2)$ are i.i.d. drawn from F_1 and F_2 separately. We can see that $h_{\mathbf{m}} \in \mathscr{L}^p(\mathbf{F}_c)$ and $h_{\mathbf{m}} \in \mathscr{O}^p(\mathbf{F}_c)$ implies $h_{\mathbf{m}} \in L^p(\mathbf{F}_c)$ and $h_{\mathbf{m}} \in O^p(\mathbf{F}_c)$ respectively, as both require the expected value remains finite regardless of repetitions among the indices.

3.2 Asymptotics of Multisample U-Statistics

We are now ready to introduce the Law of Large Numbers (LLN) for the multisample U-statistics and V-statistics. For brevity, we denote $\mu = \mathbb{E}|h_{\mathbf{m}}(\xi^{(1)};\xi^{(2)};\ldots;\xi^{(k)})|$.

Theorem 1 (Multisample Law of Large Numbers) With Definitions 2, 3 and 4, we have the following:

(a) $h_{\mathbf{m}} \in L^1(\mathbf{F}_c)$ implies $U_{\mathbf{N}}[h_{\mathbf{m}}] \xrightarrow{L^1} \mu$.

(b) $h_{\mathbf{m}} \in O^1(\mathbf{F}_c)$ implies $U_{\mathbf{N}}[h_{\mathbf{m}}] \xrightarrow{a.s.} \mu$.

(c) $h_{\mathbf{m}} \in \mathscr{L}^1(\mathbf{F}_c)$ implies $V_{\mathbf{N}}[h_{\mathbf{m}}] \xrightarrow{L^1} \mu$.

(d) $h_{\mathbf{m}} \in \mathscr{O}^1(\mathbf{F}_c)$ implies $V_{\mathbf{N}}[h_{\mathbf{m}}] \xrightarrow{a.s.} \mu$.

Sen (1977) first proved (b) and (d). The multi-index backward martingale theory developed in Gut (1976) almost implies all of (a) (b) (c) and (d) as special cases. However, there are two missing arguments: an explicit backward martingale construction (Christofides and Serfling 1990) and a multisample Hewitt Savage 0-1 law, both of which can be established with some efforts.

Unlike the LLNs, we must impose an index sequence to derive the multisample Central Limit Theorem (CLT), which reduces the asymptotic analysis from an array to a sequence. An index sequence is denoted by $\{\mathbf{N}_\ell\}_{\ell=0}^{\infty}$, where each $\mathbf{N}_\ell$ is a c-dimensional integer vector. This sequence satisfies the following conditions: (i) $\mathbf{N}_0 = (0,0,\ldots,0)$; (ii) $\mathbf{N}_{\ell+1} \geq \mathbf{N}_\ell$ according to the elementwise partial order; (iii) $\sum_{k=1}^{c} N_k \equiv \ell$. Then, a stochastic array $S_{\mathbf{N}}$ on a given index sequence $\{\mathbf{N}_\ell\}_{\ell=0}^{\infty}$ is a random sequence indexed by the norm $|\mathbf{N}| \coloneqq \sum_{k=1}^{c} N_k$. To establish a multisample CLT, we make following assumption:

Assumption 1 The sample size vector $\mathbf{N}$ increases according to a given index sequence on which the ratio $N_k/|\mathbf{N}|$ converges to $\gamma_k \in (0,1)$. That is, $\lim_{|\mathbf{N}|\to\infty} N_k/|\mathbf{N}| = \gamma_k > 0$ and $\sum_{k=1}^{c} \gamma_k = 1$.

To see the necessity of such an assumption, consider $S(N_1,N_2) = \sum_{i=1}^{N_1} X_i + \sum_{j=1}^{N_2} Y_j$ where $\{X_i\}_{i=1}^{\infty}$ and $\{Y_j\}_{j=1}^{\infty}$ are i.i.d. zero-mean Gaussian variables with *distinct* variances σ_X^2 and σ_Y^2. For a zigzag index sequence such that the ratio N_1/N_2 diverges and oscillates between two values, the variance of $S(N_1,N_2)$ will also oscillate. Consequently, even when $S_{\mathbf{N}}$ is properly scaled, no weak convergence can be concluded.

Theorem 2 (Multisample CLT) Suppose Assumption 1 holds and assume the kernel $h_{\mathbf{m}} \in L^2(\mathbf{F}_c)$. Then

$$\lim_{|\mathbf{N}|\to\infty} \sqrt{|\mathbf{N}|} \times (U_{\mathbf{N}}[h_{\mathbf{m}}] - \mu) \xrightarrow{d} \mathcal{N}\left(0, \sum_{k=1}^{c} \frac{m_k^2}{\gamma_k} \mathbb{V}\mathrm{ar}\left(\mathbb{E}[h_{\mathbf{m}}(\xi^{(1)};\xi^{(2)};\ldots;\xi^{(c)})|\xi_1^{(k)}]\right)\right), \tag{4}$$

where $\mathcal{N}$ stands for the Gaussian distribution. The same convergence holds for $V_{\mathbf{N}}[h_{\mathbf{m}}]$ if $h_{\mathbf{m}} \in \mathscr{L}^2(\mathbf{F}_c)$.

The CLT for a multisample U-statistic has been proved in Lehmann (1951) and Dwass (1956).

4 SOLUTION CONSISTENCY AND VARIANCE ESTIMATION OF MSAA

We now present our main results on the statistical properties of MSAA. We first generalize the LLN for multisample U-statistics in Section 3.2 to a uniform LLN for multisample U-processes (Section 4.1). Using this, we study the consistency (Section 4.2) and associated variance estimation (Section 4.3) of MSAA.

4.1 Multisample Wald's Uniform Law of Large Numbers

We consider a parameterized family of function kernels $\{h_{\mathbf{m}}(\theta;\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(c)})\}_{\theta\in\Theta}$ and recall the mean function $\mu(\theta) \coloneqq \mathbb{E}[h_{\mathbf{m}}(\theta;\xi^{(1)},\xi^{(2)},\ldots,\xi^{(c)})]$. The parameterized $U_{\mathbf{N}}[h_{\mathbf{m}}(\theta)]$ and $V_{\mathbf{N}}[h_{\mathbf{m}}(\theta)]$ define the U-process and V-process respectively. Analogous to the classical Wald's Uniform Law of Large Numbers (ULLN), we make the following assumptions:

Assumption 2 We have the following:

(i) There exists a metric ρ such that $(\Theta;\rho)$ is a compact space.
(ii) For any fixed $(\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(c)})$, $h_{\mathbf{m}}(\theta;\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(c)})$ is a continuous function on $(\Theta;\rho)$.
(iii) There exists a dominating kernel function $H_{\mathbf{m}}$ such that for all $(\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(c)}) \in \mathbb{R}^{|\mathbf{m}|}$

$$\sup_{\theta\in\Theta} |h(\theta;\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(c)})| \le H_{\mathbf{m}}(\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(c)}).$$

With Assumption 2, a multisample ULLN holds for $H_{\mathbf{m}}$ under different additional assumptions. This is presented in the following result:

Theorem 3 (Multisample Wald's ULLNs) Recall Definition 4. Under Assumption 2,

(a) $H_{\mathbf{m}} \in L^1(\mathbf{F}_c)$ implies $\sup_{\theta\in\Theta} |U_{\mathbf{N}}[h_{\mathbf{m}}(\theta)] - \mu(\theta)| \xrightarrow{\mathbb{P}} 0$.
(b) $H_{\mathbf{m}} \in O^1(\mathbf{F}_c)$ implies $\sup_{\theta\in\Theta} |U_{\mathbf{N}}[h_{\mathbf{m}}(\theta)] - \mu(\theta)| \xrightarrow{a.s.} 0$.
(c) $H_{\mathbf{m}} \in \mathscr{L}^1(\mathbf{F}_c)$ implies $\sup_{\theta\in\Theta} |V_{\mathbf{N}}[h_{\mathbf{m}}(\theta)] - \mu(\theta)| \xrightarrow{\mathbb{P}} 0$.
(d) $H_{\mathbf{m}} \in \mathscr{O}^1(\mathbf{F}_c)$ implies $\sup_{\theta\in\Theta} |V_{\mathbf{N}}[h_{\mathbf{m}}(\theta)] - \mu(\theta)| \xrightarrow{a.s.} 0$.

We outline the main idea of proving Theorem 3. Given the ρ-continuity and the dominating function $H_{\mathbf{m}}$, we can prove the uniform convergence up to an error interval $(-\varepsilon,\varepsilon)$ on an open ball B_ε centered at any arbitrary point $\theta' \in \Theta$, provided the radius of the open ball B_ε is sufficiently small. That is, $\sup_{\theta\in B_\varepsilon} |U_{\mathbf{N}}[h_{\mathbf{m}}(\theta)] - \mu(\theta)| \le \varepsilon$ eventually, where the multisample LLNs proved in Theorem 1 are involved. The type of convergence (in probability or almost surely) and whether it extends to the V-process depends on the assumptions made on $H_{\mathbf{m}}$. Moreover, for any $\varepsilon > 0$, a finite number of such open balls can cover Θ, implied by the compactness assumption. Consequently, the convergence on small open balls extends to the entire metric space (Θ,ρ). This proof is similar to that in Wald (1949) and Chapter 2.6, Lemma A1 of Rubinstein and Shapiro (1993), but adapted to the multisample LLNs in a tedious fashion.

4.2 Solution Consistency of MSAA

We now consider the objective function $\mu(\theta)$ in the stochastic program (2) and the MSAA objective function $\mu_{\mathbf{N}}(\theta)$ in (3). The latter, $\mu_{\mathbf{N}}(\theta)$, is a multisample U-process (or V-process, which is identical in this case) characterized by $\mathbf{m} = (1,1,\ldots,1)$, with $\mu(\theta)$ representing its mean function. As a consequence, Theorem 3 implies the consistency of the MSAA objective value estimator $\mu^*_{\mathbf{N}}$. This is presented in Corollary 4. For the special case where $\mathbf{m} = (1,1,\ldots,1)$, we will omit the subscript $\mathbf{m}$ in subsequent discussions.

Corollary 4 (MSAA Consistency) Assume the function $h(\theta;\cdot)$ in (2) satisfies Assumption 2 with a dominating function $H \in L^1(\mathbf{F}_c)$. Then $|\mu^*_{\mathbf{N}} - \mu^*| \xrightarrow{\mathbb{P}} 0$. Besides, $H \in O^1(\mathbf{F}_c)$ implies $|\mu^*_{\mathbf{N}} - \mu^*| \xrightarrow{a.s.} 0$ too.

We briefly discuss the potential strengthening of our consistency results in Corollary 4. Throughout this paper, we make the two assumptions on stagewise independence and parametrized policy. The first assumption, stagewise independence, appears challenging to relax for a similar result, as the empirical estimate $\mu_{\mathbf{N}}(\theta)$ relies on the independence assumption represented by $\mathbf{F}_c = F_1 \times F_2 \times \ldots \times F_c$. However, for the second assumption, instead of the conventional Euclidean policy parametrization, Θ can be a space of decision rules. For instance, let Θ_k be a family of continuous functions mapping the realized information at time k, denoted as $(\xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(k)}) \in \mathbb{R}^{d_k}$, to a decision space $D_k \subset \mathbb{R}^{d'_k}$ (in control theory, this is called a closed-loop control). If each D_k is compact, all $\xi^{(k)}$ are bounded, and Θ_k is equicontinuous, then the composite policy space $\Theta := \prod_{k=1}^{c} \Theta_k$ is compact with respect to the supremum norm. Furthermore, if the expected cost $\mu(\theta)$ remains continuous to the control policy $\theta_k \in \Theta_k$, the multisample Wald's ULLNs can be applied to derive MSAA consistency. This can expand the scope of applicability of our approach.

4.3 Variance Estimation

In addition to solution consistency, we are interested in quantifying the uncertainty or error of the solution. To do so, we adapt the Single Replication Procedure (SRP) (Bayraksan and Morton 2006) to multistage. Achieving this requires estimating the asymptotic variance, displayed in (4), uniformly across all $\theta \in \Theta$. When $\mathbf{m} = (1,1,\ldots,1)$, we suppress the subscript $\mathbf{m}$ and consider

$$S^{(1)}_{\mathbf{N}}[h(\theta)] := \frac{1}{N_1}\sum_{i_1=1}^{N_1}\left(\left(\prod_{k=2}^{c}\frac{1}{N_k}\right)\sum_{i_2=1}^{N_2}\cdots\sum_{i_c=1}^{N_c} h(\theta;\xi^{(1)}_{i_1};\xi^{(2)}_{i_2};\ldots;\xi^{(c)}_{i_c}) - \mu_{\mathbf{N}}(\theta)\right)^2 \tag{5}$$

as an estimator of the variance of conditional expectation $s_1[h(\theta)] := \mathbb{V}\mathrm{ar}(\mathbb{E}[h(\theta;\xi^{(1)};\xi^{(2)};\ldots;\xi^{(c)})|\xi^{(1)}])$. The next corollary states that this estimator $S^{(1)}_{\mathbf{N}}[h(\theta)]$ is uniformly consistent over Θ:

Corollary 5 (Consistency of Variance Estimator) Suppose the objective function $h(\theta;\cdot)$ in (2) satisfies Assumption 2 with a dominating function $H \in L^2(\mathbf{F}_c)$. Then $\sup_{\theta\in\Theta}|S^{(1)}_{\mathbf{N}}[h(\theta)] - s_1[h(\theta)]| \xrightarrow{\mathbb{P}} 0$. Besides, if $H \in O^2(\mathbf{F}_c)$, the a.s.-convergence holds too.

Proof. Define $g(\theta;\cdot) := h(\theta;\cdot) - \mu(\theta)$. Then

$$\begin{aligned} S^{(1)}_{\mathbf{N}}[h(\theta)] &= \frac{1}{N_1}\sum_{i_1=1}^{N_1}\left(\left(\prod_{k=2}^{c}\frac{1}{N_k}\right)\sum_{i_2=1}^{N_2}\cdots\sum_{i_c=1}^{N_c} h(\theta;\xi^{(1)}_{i_1};\ldots;\xi^{(c)}_{i_c}) - \mu(\theta) - (\mu_{\mathbf{N}}(\theta)-\mu(\theta))\right)^2 \\ &= \frac{1}{N_1}\sum_{i_1=1}^{N_1}\left(\left(\prod_{k=2}^{c}\frac{1}{N_k}\right)\sum_{i_2=1}^{N_2}\cdots\sum_{i_c=1}^{N_c} g(\theta;\xi^{(1)}_{i_1};\ldots;\xi^{(c)}_{i_c})\right)^2 - (\mu_{\mathbf{N}}(\theta)-\mu(\theta))^2 \end{aligned}$$

$$= \left(\frac{1}{N_1}\prod_{k=2}^{c}\frac{1}{N_k^2}\right)\sum_{i_1=1}^{N_1}\sum_{i_2=1}^{N_2}\cdots\sum_{i_c=1}^{N_c}\sum_{j_2=1}^{N_2}\cdots\sum_{j_c=1}^{N_c} g(\theta;\xi_{i_1}^{(1)};\xi_{i_2}^{(2)};\ldots;\xi_{i_c}^{(c)})g(\theta;\xi_{i_1}^{(1)};\xi_{j_2}^{(2)};\ldots;\xi_{j_c}^{(c)}) - (\mu_{\mathbf{N}}(\theta)-\mu(\theta))^2$$

The first term can be seen as a V-process of the following kernel function with $\mathbf{m} = (1,2,2,\ldots,2)$:

$$f_{\mathbf{m}}(\theta;x^{(1)};x_1^{(2)},x_2^{(2)};x_1^{(3)},x_2^{(3)};\ldots;x_1^{(c)},x_2^{(c)}) := g(\theta;x^{(1)};x_1^{(2)};\ldots;x_1^{(c)})g(\theta;x^{(1)};x_2^{(2)};\ldots;x_2^{(c)})$$

The Cauchy–Schwarz inequality implies

$$\mathbb{E}|g(\theta;\xi_{i_1}^{(1)};\xi_{i_2}^{(2)};\ldots;\xi_{i_c}^{(c)})g(\theta;\xi_{i_1}^{(1)};\xi_{j_2}^{(2)};\ldots;\xi_{j_c}^{(c)})| \le \mathbb{E}|g(\theta;\xi_{i_1}^{(1)};\xi_{i_2}^{(2)};\ldots;\xi_{i_c}^{(c)})|^2 \le 4\mathbb{E}|H(\xi_{i_1}^{(1)};\ldots;\xi_{i_c}^{(c)})|^2 < \infty$$

The same inequality holds with arbitrary variable repetition, i.e., $i_k = j_k$ for some $k \in \{2,3,\ldots,c\}$. Denote $\nu(\theta) := \mathbb{E}[g(\theta;\xi_{i_1}^{(1)};\xi_{i_2}^{(2)};\ldots;\xi_{i_c}^{(c)})g(\theta;\xi_{i_1}^{(1)};\xi_{j_2}^{(2)};\ldots;\xi_{j_c}^{(c)})]$. Then

$$\sup_{\theta\in\Theta}\left|\left(\frac{1}{N_1}\prod_{k=2}^{c}\frac{1}{N_k^2}\right)\sum_{i_1=1}^{N_1}\sum_{i_2=1}^{N_2}\cdots\sum_{i_c=1}^{N_c}\sum_{j_2=1}^{N_2}\cdots\sum_{j_c=1}^{N_c} f_{\mathbf{m}}(\theta;\xi_{i_1}^{(1)};\xi_{i_2}^{(2)},\xi_{j_2}^{(2)};\ldots;\xi_{i_c}^{(c)},\xi_{j_c}^{(c)}) - \nu(\theta)\right| \xrightarrow{\mathbb{P}} 0, \quad (6)$$

despite that $f_{\mathbf{m}}$ may not be a symmetric function. In fact, Theorem 3 holds without the symmetric kernel assumption. According to the tower rule,

$$\begin{aligned}\nu(\theta) =& \mathbb{E}[\mathbb{E}[g(\theta;\xi_{i_1}^{(1)};\xi_{i_2}^{(2)};\ldots;\xi_{i_c}^{(c)})g(\theta;\xi_{i_1}^{(1)};\xi_{j_2}^{(2)};\ldots;\xi_{j_c}^{(c)})|\xi_{i_1}^{(1)}]]\\ =& \mathbb{E}[\Big(\mathbb{E}[g(\theta;\xi_{i_1}^{(1)};\xi_{i_2}^{(2)};\ldots;\xi_{i_c}^{(c)})|\xi_{i_1}^{(1)}]\Big)^2] = \mathbb{V}\mathrm{ar}(\mathbb{E}[h(\theta;\xi^{(1)};\xi^{(2)};\ldots;\xi^{(c)})|\xi^{(1)}]) = s_1[h(\theta)].\end{aligned}$$

Consequently,

$$\begin{aligned}&\sup_{\theta\in\Theta}|S_{\mathbf{N}}^{(1)}[h(\theta)] - s_1[h(\theta)]| \le \sup_{\theta\in\Theta}|\mu_{\mathbf{N}}(\theta)-\mu(\theta)|^2\\ &+\sup_{\theta\in\Theta}\left|\left(\frac{1}{N_1}\prod_{k=2}^{c}\frac{1}{N_k^2}\right)\sum_{i_1=1}^{N_1}\sum_{i_2=1}^{N_2}\cdots\sum_{i_c=1}^{N_c}\sum_{j_2=1}^{N_2}\cdots\sum_{j_c=1}^{N_c} f_{\mathbf{m}}(\theta;\xi_{i_1}^{(1)};\xi_{i_2}^{(2)},\xi_{j_2}^{(2)};\ldots;\xi_{i_c}^{(c)},\xi_{j_c}^{(c)}) - s_1[h(\theta)]\right|,\end{aligned}$$

where the first term converges to 0 almost surely and the second term converges to 0 in probability. Hence $\sup_{\theta\in\Theta}|S_{\mathbf{N}}^{(1)}[h_{\mathbf{m}}(\theta)] - s_1(\theta)| \xrightarrow{\mathbb{P}} 0$. The proof for a.s.-convergence is similar, where the stronger assumption $H_{\mathbf{m}} \in O^2(\mathbf{F}_c)$ implies a.s.-convergence in (6). □

5 MULTISAMPLE SINGLE REPLICATION PROCEDURE

We now consider uncertainty quantification, more concretely constructing an upper confidence bound for the optimality gap, for a given solution $\hat{\theta}$ of the stochastic program (2). More precisely, we define the optimality gap of solution $\hat{\theta}$ as $G(\hat{\theta}) := \mu(\hat{\theta}) - \mu^*$. Let $\mu_{\mathbf{N}}^*$ be the optimal value from the MSAA problem (3) with random samples $\{\xi_{i_k}^{(k)}\}_{i_k=1}^{N_k}$ for $k = 1,\ldots,c$. Then

$$\mathbb{E}[\mu_{\mathbf{N}}^*] = \mathbb{E}\left[\min_{\theta\in\Theta}\left(\prod_{k=1}^{c}\frac{1}{N_k}\right)\sum_{i_1=1}^{N_1}\sum_{i_2=1}^{N_2}\cdots\sum_{i_c=1}^{N_c} h(\theta;\xi_{i_1}^{(1)},\xi_{i_2}^{(2)},\ldots,\xi_{i_c}^{(c)})\right] \le \min_{\theta\in\Theta}\mu(\theta) = \mu^*,$$

which indicates that $\mu^*_{\mathbf{N}}$ is a lower bound of μ^* in expectation. By the definition of optimality, we have $\mathbb{E}[\mu^*_{\mathbf{N}}] \le \mu^* \le \mathbb{E}[\mu_{\mathbf{N}}(\hat{\theta})] = \mu(\hat{\theta})$ so that $\mu(\hat{\theta}) - \mathbb{E}[\mu^*_{\mathbf{N}}] \ge G(\hat{\theta})$ is an upper bound for the optimality gap. Intuitively, when $\mathbf{N} \to \infty$ elementwise, both $\mu^*_{\mathbf{N}}$ and $\mu_{\mathbf{N}}(\hat{\theta})$ would serve as good estimators for their respective mean, implied by their consistency. Hence, we estimate the optimality gap as $G_{\mathbf{N}}(\hat{\theta}) := \mu_{\mathbf{N}}(\hat{\theta}) - \mu^*_{\mathbf{N}}$, i.e.

$$G_{\mathbf{N}}(\hat{\theta}) = \left(\prod_{k=1}^{c} \frac{1}{N_k}\right) \sum_{i_1=1}^{N_1} \cdots \sum_{i_c=1}^{N_c} h(\hat{\theta}; \xi^{(1)}_{i_1}, \ldots, \xi^{(c)}_{i_c}) - \left(\prod_{k=1}^{c} \frac{1}{N_k}\right) \sum_{i_1=1}^{N_1} \cdots \sum_{i_c=1}^{N_c} h(\theta^*_{\mathbf{N}}; \xi^{(1)}_{i_1}, \ldots, \xi^{(c)}_{i_c}). \quad (7)$$

Now consider the function $f(\theta; x^{(1)}, x^{(2)}, \ldots, x^{(c)}) := h(\hat{\theta}; x^{(1)}, x^{(2)}, \ldots, x^{(c)}) - h(\theta; x^{(1)}, x^{(2)}, \ldots, x^{(c)})$ and define its mean process to be $\mu^f_{\mathbf{N}}(\theta) := \prod_{k=1}^{c} \frac{1}{N_k} \sum_{i_2=1}^{N_2} \cdots \sum_{i_c=1}^{N_c} f(\theta; \xi^{(1)}_{i_1}; \xi^{(2)}_{i_2}; \ldots; \xi^{(c)}_{i_c})$. Then the sample variance process of the ℓ-th conditional expectation is defined as

$$S^{(\ell)}_{\mathbf{N}}[f(\theta)] := \frac{1}{N_\ell} \sum_{i_\ell=1}^{N_\ell} \left(\prod_{k \ne \ell}^{c} \left(\frac{1}{N_k} \sum_{i_k=1}^{N_k} \right) \circ f(\theta; \xi^{(1)}_{i_1}; \xi^{(2)}_{i_2}; \ldots; \xi^{(c)}_{i_c}) - \mu^f_{\mathbf{N}}(\theta) \right)^2. \quad (8)$$

For a fixed θ, equation (8) serves as a natural estimator for the conditional variance. Specifically, by setting $c = 2$, suppressing θ and representing $(\xi^{(1)}_{i_1}, \xi^{(2)}_{i_2})$ as (X_i, Y_j), it simplifies to

$$S^{(1)}_{\mathbf{N}}[f] := \frac{1}{N_1} \sum_{i=1}^{N_1} \left(\sum_{j=1}^{N_2} \frac{f(X_i, Y_j) - \mu^f_{\mathbf{N}}}{N_2} \right)^2 \text{ and } S^{(2)}_{\mathbf{N}}[f] := \frac{1}{N_2} \sum_{j=1}^{N_2} \left(\sum_{i=1}^{N_1} \frac{f(X_i, Y_j) - \mu^f_{\mathbf{N}}}{N_1} \right)^2,$$

which are intuitive choices for estimating $\mathbb{V}\mathrm{ar}(\mathbb{E}[f(X_1, Y_1)|X_1])$ and $\mathbb{V}\mathrm{ar}(\mathbb{E}[f(X_1, Y_1)|Y_1])$. Now we impose Assumption 1. Recall the ratio $\gamma_k = \lim_{|\mathbf{N}| \to \infty} N_k / |\mathbf{N}|$ and define the variance process as $S_{\mathbf{N}}[f(\theta)] = \sum_{\ell=1}^{c} \frac{|\mathbf{N}|}{N_k} S^{(\ell)}_{\mathbf{N}}[f(\theta)]$. Also, recall that $s_k[f(\theta)] = \mathbb{V}\mathrm{ar}(\mathbb{E}[f(\theta; \xi^{(1)}; \xi^{(2)}; \ldots; \xi^{(c)})|\xi^{(k)}])$ and define $s[f(\theta)] = \sum_{\ell=1}^{c} \frac{1}{\gamma_\ell} s_\ell[f(\theta)]$. Then according to Corollary 5, if $\sup_{\theta \in \Theta} |f(\theta; x^{(1)}, \ldots, x^{(c)})| \in O^2(\mathbf{F}_c)$, we have the a.s.-covergence that $\sup_{\theta \in \Theta} |S_{\mathbf{N}}[f(\theta)] - s[f(\theta)]| \xrightarrow{a.s.} 0$. The analysis is formally summarized as follows:

Lemma 6 Let Θ^* be the optimal solution set of the stochastic program (2), and $\theta^*_{\min}, \theta^*_{\max} \in \Theta$ such that $s[f(\theta^*_{\min})] \le s[f(\theta^*)] \le s[f(\theta^*_{\max})]$ for all $\theta^* \in \Theta$. Under Assumptions 1 and 2 and that the dominating function $H \in O^2(\mathbf{F}_c)$, we assert (a) $\mu^*_{\mathbf{N}} \xrightarrow{a.s.} \mu^*$; (b) all limit points of $\theta^*_{\mathbf{N}}$ lie in Θ^*; (c) with probability 1 that

$$s[f(\theta^*_{\min})] \le \liminf_{|\mathbf{N}| \to \infty} S_{\mathbf{N}}[f(\theta^*_{\mathbf{N}})] \le \limsup_{|\mathbf{N}| \to \infty} S_{\mathbf{N}}[f(\theta^*_{\mathbf{N}})] \le s[f(\theta^*_{\max})].$$

In Lemma 6, both (a) and (b) are trivially implied by the uniform convergence established in Corollary 5. Due to the compactness of Θ, the sequence $\{\theta^*_{\mathbf{N}}\}_{|\mathbf{N}|=1}^{\infty}$ can be decomposed into at most countable convergent subsequences. For each subsequence, $s[f(\theta^*_{\min})] \le \lim_{|\mathbf{N}| \to \infty} S_{\mathbf{N}}[f(\theta^*_{\mathbf{N}})] \le s[f(\theta^*_{\max})]$ with probability 1. Hence (c) follows.

Based on Lemma 6, we propose the Multisample Single Replication Procedure (MSRP) that generalizes SRP in a natural fashion. This is presented in Algorithm 1, which exhibits asymptotically valid coverage as described in the theorem below. The notation $\Phi^{-1}_{1-\alpha}$ stands for the $(1-\alpha)$-quantile of a standard normal distribution. Its proof is similar to that used for SRP in Bayraksan and Morton (2006). Intuitively, a

uniform convergence estimator for the asymptotic variance is needed for this procedure, as it utilizes the same dataset to solve the MSAA for $\theta^*_{\mathbf{N}}$ and to estimate the asymptotic variance at $\theta^*_{\mathbf{N}}$ as specified by the CLT in Theorem 2. Relying on uniformity conveniently circumvent potential issues with correlation. These intuitions follow the elegant arguments of Bayraksan and Morton (2006).

Algorithm 1 Multisample Single Replication Procedure

1: **procedure** INPUT(confidence level $\alpha \in (0,1)$; candidate policy $\hat{\theta} \in \Theta$)
2: Samples $\{\xi_{i_k}^{(k)}\}_{i_k=1}^{N_k}$ with sizes $\mathbf{N} = (N_1, N_2, \ldots, N_c)$
3: Minimize $\mu_{\mathbf{N}}(\theta)$ over Θ to obtain solution $\theta^*_{\mathbf{N}}$.
4: Calculate $G_{\mathbf{N}}(\hat{\theta})$ as defined in (7) and $S^{(\ell)}_{\mathbf{N}}[f(\theta^*_{\mathbf{N}})]$ as defined in (8).
5: **return** upper confidence bound $G_{\mathbf{N}}(\hat{\theta}) + \sqrt{\sum_{\ell=1}^{c} \frac{1}{N_\ell} S^{(\ell)}_{\mathbf{N}}[f(\theta^*_{\mathbf{N}})]} \times \Phi^{-1}_{1-\alpha}$
6: **end procedure**

Theorem 7 (Asymptotically Valid Coverage) Suppose Assumptions 1 and 2 hold and that the dominating function $H^2 \in O(\mathbf{F}_c)$. Given $\alpha \in (0, \frac{1}{2})$ and letting $\Phi^{-1}_{1-\alpha}$ be the $(1-\alpha)$-quantile of the standard normal distribution,

$$\liminf_{|\mathbf{N}|\to\infty} \mathbb{P}\left(G(\hat{\theta}) \le G_{\mathbf{N}}(\hat{\theta}) + \sqrt{\sum_{\ell=1}^{c} \frac{1}{N_\ell} S^{(\ell)}_{\mathbf{N}}[f(\theta^*_{\mathbf{N}})]} \times \Phi^{-1}_{1-\alpha} \right) \ge 1-\alpha.$$

ACKNOWLEDGMENTS

We gratefully acknowledge support from the InnoHK initiative, the Government of the HKSAR, Laboratory for AI-Powered Financial Technologies, and the Columbia SEAS Innovation Hub grant.

REFERENCES

Arcones, M. A. and E. Giné. 1993. "Limit theorems for U-processes". *The Annals of Probability*:1494–1542.

Bayraksan, G. and D. P. Morton. 2006. "Assessing solution quality in stochastic programs". *Mathematical Programming* 108:495–514.

Chen, X. and D. L. Woodruff. 2023. "Software for data-based stochastic programming using bootstrap estimation". *INFORMS Journal on Computing* 35(6):1218–1224.

Christofides, T. C. and R. J. Serfling. 1990. "Maximal inequalities and convergence results for generalized U-statistics". *Journal of statistical planning and inference* 24(3):271–286.

de la Pena, V. H. 1992. "Decoupling and Khintchine's inequalities for U-statistics". *The Annals of Probability*:1877–1892.

Dwass, M. 1956. "The large-sample power of rank order tests in the two-sample problem". *The Annals of Mathematical Statistics*:352–374.

Eichhorn, A. and W. Römisch. 2007. "Stochastic integer programming: Limit theorems and confidence intervals". *Mathematics of Operations Research* 32(1):118–135.

Gut, A. 1976. "Convergence of reversed martingales with multidimensional indices". *Duke Math Journal* 43(2).

Koussoulas, N. T. and C. T. Leondes. 1986. "The multiple linear quadratic Gaussian problem". *International Journal of Control* 43(2):337–349.

Lam, H. and H. Qian. 2018a. "Assessing solution quality in stochastic optimization via bootstrap aggregating". In *2018 Winter Simulation Conference (WSC)*, 2061–2071. IEEE.

Lam, H. and H. Qian. 2018b. "Bounding optimality gap in stochastic optimization via bagging: Statistical efficiency and stability". *arXiv preprint arXiv:1810.02905*.

Lehmann, E. L. 1951. "Consistency and unbiasedness of certain nonparametric tests". *The annals of mathematical statistics*:165–179.

Mak, W.-K., D. P. Morton, and R. K. Wood. 1999. "Monte Carlo bounding techniques for determining solution quality in stochastic programs". *Operations research letters* 24(1-2):47–56.

Moore, J. B., X. Y. Zhou, and A. E. Lim. 1999. "Discrete time LQG controls with control dependent noise". *Systems & Control Letters* 36(3):199–206.

Rubinstein, R. Y. and A. Shapiro. 1993. *Discrete event systems: sensitivity analysis and stochastic optimization by the score function method*, Volume 13. Wiley.

Scarf 1960. "The Optimality of (s, S) Policies in the Dynamic Inventory Problem". *Mathematical Methods in the Social Sciences*.

Sen, P. K. 1974. "Weak convergence of generalized U-statistics". *The Annals of Probability*:90–102.

Sen, P. K. 1977. "Almost sure convergence of generalized U-statistics". *The Annals of Probability*:287–290.

Shapiro, A. 2003. "Inference of statistical bounds for multistage stochastic programming problems". *Mathematical Methods of Operations Research* 58(1):57–68.

Shapiro, A. 2008. "Stochastic programming approach to optimization under uncertainty". *Mathematical Programming* 112:183–220.

Shapiro, A., D. Dentcheva, and A. Ruszczynski. 2021. *Lectures on stochastic programming: modeling and theory*. SIAM.

Shapiro, A. and A. Nemirovski. 2005. "On complexity of stochastic programming problems". *Continuous optimization: Current trends and modern applications*:111–146.

Smythe, R. T. 1973. "Strong laws of large numbers for r-dimensional arrays of random variables". *The Annals of Probability*:164–170.

Swamy, C. and D. B. Shmoys. 2005. "Sampling-based approximation algorithms for multi-stage stochastic optimization". In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, 357–366. IEEE.

Talluri, K. T. and G. J. Van Ryzin. 2006. *The theory and practice of revenue management*, Volume 68. Springer Science & Business Media.

Wald, A. 1949. "Note on the consistency of the maximum likelihood estimate". *The Annals of Mathematical Statistics* 20(4):595–601.

AUTHOR BIOGRAPHIES

YUNHAO YAN is a Ph.D. student in the Department of Industrial Engineering and Operations Research at Columbia University. His email address is yy2882@columbia.edu.

HENRY LAM is an Associate Professor in the Department of Industrial Engineering and Operations Research at Columbia University. His email address is henry.lam@columbia.edu.