

## OPTIMAL ALLOCATIONS FOR SAMPLE AVERAGE APPROXIMATION

Prateek Jaiswal  
Harsha Honnappa

School of Industrial Engineering  
Purdue University  
West Lafayette, IN 47906, USA

Raghu Pasupathy

Department of Statistics  
Purdue University  
West Lafayette, IN 47907, USA

### ABSTRACT

We consider a single stage stochastic program without recourse with a strictly convex loss function. We assume a compact decision space and grid it with a finite set of points. In addition, we assume that the decision maker can generate samples of the stochastic variable independently at each grid point and form a *sample average approximation* (SAA) of the stochastic program. Our objective in this paper is to characterize an asymptotically optimal linear sample allocation rule, given a fixed sampling budget, which maximizes the decay rate of probability of making false decision.

### 1 INTRODUCTION

Let  $\xi$  be a measurable function that induces the distribution function  $F(x) := \mathbb{P}(\xi \leq x)$ . Stochastic programs are canonical models for decision-making under uncertainty, covering a multitude of classic stochastic optimization problems:

$$\begin{aligned} & \text{minimize} && f(x) := \mathbb{E}[L(x, \xi)] = \int_{\Xi} L(x, \xi) dF(\xi), \\ & \text{subject to} && x \in \mathcal{X}, \end{aligned} \tag{1}$$

where  $L(x, \xi) : \mathcal{X} \times \Xi \rightarrow \mathbb{R}$  is a continuously differentiable ‘loss’ function,  $\mathcal{X} \subset \mathbb{R}$  is the set of decision variables; for brevity we also call the sample values as  $\xi$ . In this paper, we assume that  $L(x, \xi)$  is convex in  $x$  and measurable with respect to  $\xi$ .

*Sample average approximation* (SAA) is a classic Monte Carlo method for estimating the stochastic program (1). Here, the decision maker (DM) grids the decision space  $\mathcal{X}$  into a finite set of points  $D := \{x_1, x_2, \dots, x_d\}$ , and simulates samples of the loss function  $L(x_i, \xi)$  at each of the grid points. Assuming a total sampling budget of  $n$ , the DM generates  $m(x_i)$  independent and identically distributed (i.i.d) samples of  $L(x_i, \xi)$ , denoted as  $L(x_i, \xi_j^i) \forall 1 \leq j \leq m(x_i)$ ; at each point  $x_i$  with  $\sum_{i=1}^d m(x_i) = n$ . Then, the SAA stochastic program is

$$\min_{x_i \in D} \hat{f}(x_i) := \frac{1}{m(x_i)} \sum_{j=1}^{m(x_i)} L(x_i, \xi_j^i). \tag{2}$$

Modulo regularity conditions on  $L(\cdot, \cdot)$ , the strong law of large numbers (SLLN) implies that (2) converges to the true program in (1) (Shapiro, Dentcheva, and Ruszczyński 2009). It is also known that the optimizers are consistent and the optimal rate of convergence is  $O(n^{-1/2})$ . An important allied question to these asymptotic results is a quantification of the likelihood that the empirical optimizer diverges from the true optimizer for a given sampling budget  $n$ . We seek such a quantification for two reasons:

- A quantification of this rate gives a clear sense of how ‘good’ the empirical optimizer and empirical optimal value are, and

- in the simulation context it provides a guideline on how to allocate a limited sampling/computational budget across the design points in  $D$ .

In this paper, we focus on the latter issue. In general, it is a formidable task to compute the likelihood for a fixed budget; in full generality, one requires tight concentration bounds in order to make meaningful predictions about budget allocations. Instead, in this paper we establish a large deviations principle (LDP) satisfied by the Monte Carlo estimator (2) as the sampling budget tends to infinity. For the definition of LDP, we refer the readers to section 1.2 of Dembo and Zeitouni (2010).

It is important here to differentiate between the optimally computing budget allocation (OCBA) method for selecting an optimal system from a finite set of systems (Chen et al. 2000; Glynn and Juneja 2004) and our approach to SAA problem. In OCBA, there is no topology associated with the finite set of systems unlike SAA. In addition, OCBA approach only considers probability of selecting suboptimal system due to random sampling errors, whereas our framework also takes into account the discretization error.

### 1.1 Our Contributions

Our main objective is the derivation of an optimal allocation of the sampling budget across the design points such that a canonical LD rate is achieved at the optimizer of (2). In particular, we seek what we term as ‘linear’ allocation rules where  $m(x) = \alpha_x n$ , where  $\alpha_x \in [0, 1]$  and  $\sum_{x \in D} \alpha_x = 1$ .

We make the simplifying assumption that the DM can sample independently from each design point  $x \in D$ . In effect, this allows an ‘embarrassingly’ parallel implementation of the SAA estimation, where ‘slave’ machines compute  $\hat{f}(x)$  with  $m(x)$  samples, and communicate the result to a central ‘master’ machine that coordinates the budget allocation and aggregates the calculations to compute (2).

Now let  $\hat{x}_n := \arg \min_{x_i \in D} \hat{f}(x_i)$ ,  $\hat{x} := \arg \min_{x_i \in D} f(x_i)$ , and  $x^* := \arg \min_{x \in \mathcal{X}} f(x)$ . As noted before, the Monte Carlo empirical objective converges to the population objective as  $n \rightarrow \infty$  for any linear allocation rule, and consequently one expects that  $\hat{x}_n \rightarrow \hat{x}$  almost surely (a.s.) as  $n \rightarrow \infty$ . Since the objective is assumed to be continuously differentiable, it follows that  $f(\hat{x}_n) \rightarrow f(\hat{x})$  a.s. as  $n \rightarrow \infty$ . With this information, we establish the canonical LD rate function satisfied by the optimal value of (2) as a function of some linear allocation rule and in the limit of a large sampling budget. This result follows from the Gärtner-Ellis theorem, and does not involve any analytical subtleties in light of our assumptions. However, the rate function has not appeared in the literature before, and might be of interest more generally. The proof proceeds in two steps. We first characterize the LD rate of the likelihood of mis-ordering the SAA empirical estimates at any two points in the decision space  $D$ . Next, we use this result to establish our main result on the LD rate on the likelihood that the objective value  $f(\hat{x}_n)$  at the SAA optimizer  $\hat{x}_n$  in (2) is at least  $\varepsilon > 0$  worse than the true value  $f(x^*)$ , in the large budget limit.

We next provide structural results on the LD rate function, in particular demonstrating that it is strictly concave in the allocation rule  $\alpha := (\alpha_x, x \in D)$ . Consequently, there exists a unique optimal linear allocation rule. This result, of course, presumes that the master machine has complete information about the statistics of the expected loss function - in particular, we assume the existence of a cumulant generating function. In practice, this is not an implementable policy, since the DM only has access to a Monte Carlo simulator. We next design two recursive algorithms that optimize the LD rate function as samples accumulate. The first algorithm parallels Algorithm 2 in Hunter and Pasupathy (2013) and is applicable when a closed form expression for the rate function for the mis-ordering likelihood is available. When such an expression is available the problem is really one of ranking and selection (R&S). In general SAA problems, closed-forms are not easy to compute and this too must be estimated. Our second algorithm is an ‘expectation-maximization’ style recursive algorithm. We illustrate these algorithms with numerical simulation results.

The remainder of the paper is organized as follows. We begin in Section 3 by proving the LDP satisfied by the SAA estimator. In Section 4 we derive structural properties of the LD rate function as function of the linear allocation rule, and exhibit the variational optimization problem to find the optimal linear

allocation rule. We then derive a recursive algorithm for computing the optimal allocation rule on a sample path (and fixed sampling budget), and illustrate the algorithm, on three different example problems. We end with comments on several future directions for this paper.

## 2 NOTATIONS AND PRELIMINARIES

We assume there exists a probability sample space  $(\Omega, \mathfrak{E}, \mathbb{P})$ , and define  $\xi$  with respect to this space. The indicator function of a set is represented by  $\mathbb{I}_{\{\cdot\}}$  and  $\lfloor \cdot \rfloor$  denotes the greatest integer function. We define a ‘regret’ function

$$[f(\hat{x}_n) - f(x^*)] = [f(\hat{x}_n) - f(\hat{x})] + [f(\hat{x}) - f(x^*)], \quad (3)$$

where the first term on the right hand side is the *Sampling Error* and the latter term the *Discretization Error*. We make a few assumptions to guarantee the existence of the LD rate function. We assume that the loss function  $L(\cdot, \cdot)$  satisfies

**Assumption 1**  $L(x, \xi)$  is not a point mass at  $f(x)$  for all  $x \in \mathcal{X}$  and for some continuously differentiable function  $f$ .

**Assumption 2** The cumulant generating function (CGF) of  $L(\cdot, \xi)$  is well defined and finite for all  $x \in \mathcal{X}$ , that is

$$\Lambda(x, \theta) := \log \mathbb{E} \left[ e^{\theta L(x, \xi)} \right] < \infty \quad \forall \theta \in \mathbb{R}, \quad \forall x \in \mathcal{X}.$$

When the loss function  $L(\cdot, \xi)$  is bounded above by  $\xi$  and the CGF of  $\xi$  is well defined and finite, the above assumption is trivially satisfied.

**Assumption 3** Let  $H_x(\Lambda) := \{\theta : \Lambda(x, \theta) < \infty\} \forall x \in \mathcal{X}$  be such that the origin belongs to the interior of  $H_x(\Lambda)$ . Furthermore, we also assume that  $\Lambda(x, \theta)$  is steep, that is  $\lim_{\theta \rightarrow \infty} \left| \frac{\partial \Lambda(x, \theta)}{\partial \theta} \right|_{\theta = \theta_n} = \infty$ , for any sequence  $\theta_n$  in the interior of  $H_x(\Lambda)$ , which converges to a boundary point of  $H_x(\Lambda) \forall x \in \mathcal{X}$ .

## 3 LARGE DEVIATIONS FOR SAA

In this section we establish an LDP satisfied by the regret function. Let  $x, y \in D$  and consider the Monte Carlo estimates  $\hat{f}(x)$  and  $\hat{f}(y)$ . Our first result establishes a LD rate function for the likelihood that  $\hat{f}(x)$  and  $\hat{f}(y)$  are mis-ordered in the large budget limit.

**Lemma 1** Fix  $\gamma > 0$  and  $x, y \in D$ . Then, under Assumptions 1, 2 and 3,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\hat{f}(y) - \hat{f}(x) \geq \gamma) = -I(\gamma, \alpha_x, \alpha_y) \mathbb{I}_{\{f(y) - f(x) < \gamma\}},$$

where  $I(\gamma, \alpha_x, \alpha_y) := \sup_{t \in \mathbb{R}} \left( t\gamma - \alpha_y \Lambda \left( y, \frac{t}{\alpha_y} \right) - \alpha_x \Lambda \left( x, -\frac{t}{\alpha_x} \right) \right)$ .

*Proof.* Let  $Y_n := n(\hat{f}(y) - \hat{f}(x))$ , and observe that for any  $t \in \mathbb{R}$

$$\begin{aligned} \mathbb{E}[e^{tY_n}] &= \mathbb{E} \left[ \exp \left( t \sum_{i=1}^{\lfloor \alpha_y n \rfloor} \frac{L(y, \xi_i^y)}{\alpha_y} - t \sum_{j=1}^{\lfloor \alpha_x n \rfloor} \frac{L(x, \xi_j^x)}{\alpha_x} \right) \right] \\ &= \mathbb{E} \left[ \exp \left( t \frac{L(y, \xi)}{\alpha_y} \right) \right]^{\lfloor \alpha_y n \rfloor} \mathbb{E} \left[ \exp \left( -t \frac{L(x, \xi)}{\alpha_x} \right) \right]^{\lfloor \alpha_x n \rfloor}, \end{aligned}$$

where the last equality follows from the fact that we sample independently at every design point. Next, using the fact that  $\lim_{n \rightarrow \infty} \lfloor \alpha_i n \rfloor / n = \alpha_i$  observe that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} [e^{tY_n}] = \alpha_y \log \mathbb{E} \left[ \exp \left( t \frac{L(y, \xi)}{\alpha_y} \right) \right] + \alpha_x \log \mathbb{E} \left[ \exp \left( -t \frac{L(x, \xi)}{\alpha_x} \right) \right] =: \varphi(t, \alpha_x, \alpha_y) \quad \forall t \in \mathbb{R}.$$

By Assumption 2,  $\varphi(t, \alpha_x, \alpha_y) < \infty \quad \forall t \in \mathbb{R}$ . Together with Assumption 3 it follows that the Gärtner–Ellis Theorem holds (Dembo and Zeitouni 2010), and the lemma is proved with good rate function  $I(z, \alpha_x, \alpha_y) := \sup_{t \in \mathbb{R}} (tz - \varphi(t, \alpha_x, \alpha_y))$ . Since  $I(z, \alpha_x, \alpha_y)$  is strictly convex in  $z$  and attains the minimum value 0 precisely at  $f(y) - f(x)$ , Therefore,  $\inf_{z \in [\gamma, \infty)} I(z, \alpha_x, \alpha_y) = I(\gamma, \alpha_x, \alpha_y) \mathbb{I}_{\{f(y) - f(x) < \gamma\}}$ .  $\square$

The next two lemmas are crucial for establishing the main result of this section.

**Lemma 2** Fix  $k \in \mathbb{N}$  and let  $\{a_i^n\} \subset \mathbb{R}$  be arbitrary sequences for  $1 \leq i \leq k$ . Then

$$\liminf_{n \rightarrow \infty} \max\{a_1^n, a_2^n, \dots, a_k^n\} \geq \max\{\liminf_{n \rightarrow \infty} a_1^n, \liminf_{n \rightarrow \infty} a_2^n, \dots, \liminf_{n \rightarrow \infty} a_k^n\}.$$

*Proof.* First observe that for any  $i \in \{1, 2, \dots, k\}$   $\max\{a_1^n, a_2^n, \dots, a_k^n\} \geq a_i^n$ . Therefore,  $\liminf_{n \rightarrow \infty} \max\{a_1^n, a_2^n, \dots, a_k^n\} \geq \liminf_{n \rightarrow \infty} a_i^n$ . Since this holds for any  $i \in \{1, 2, \dots, k\}$ , the lemma follows.  $\square$

**Lemma 3** Let  $\{a_i^n\} \subset \mathbb{R}$  for  $i = 1, 2$  be arbitrary sequences. Then,

$$\liminf_{n \rightarrow \infty} (a_1^n + a_2^n) \geq \liminf_{n \rightarrow \infty} a_1^n + \liminf_{n \rightarrow \infty} a_2^n.$$

*Proof.* The proof follows from the definition of  $\liminf$  and infimum inequality.  $\square$

We now turn to main result, which establishes an LDP for the regret (3).

**Theorem 1** Fix  $\varepsilon > 0$ . Under Assumptions 1, 2 and 3 the regret (3) satisfies

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(f(\hat{x}_n) - f(x^*) \geq \varepsilon) = -J(\varepsilon),$$

where  $J(\varepsilon) := \min_{x \in Q(\delta)} \sum_{y \in D} I(\alpha_x, \alpha_y) \mathbb{I}_{\{f(y) < f(x)\}}$ ,  $I(\alpha_x, \alpha_y) := \sup_{t \in \mathbb{R}} \left( -\alpha_y \Lambda(y, \frac{t}{\alpha_y}) - \alpha_x \Lambda(x, -\frac{t}{\alpha_x}) \right)$ , and  $Q(\delta) := \{x \in D : f(x) > f(x^*) + \varepsilon\}$ .

*Proof.* Recall that our objective is to characterize the rate of decay of the likelihood of the rare event,

$$\{f(\hat{x}_n) - f(x^*) > \varepsilon\}. \tag{4}$$

Observe that

$$\begin{aligned} \mathbb{P}(f(\hat{x}_n) - f(x^*) > \varepsilon) &= \mathbb{P}(f(\hat{x}_n) - f(\hat{x}) + f(\hat{x}) - f(x^*) > \varepsilon) \\ &= \mathbb{P}(f(\hat{x}_n) - f(\hat{x}) > \delta), \end{aligned}$$

where  $\delta := \varepsilon - (f(\hat{x}) - f(x^*))$ , and  $f(\hat{x}) - f(x^*)$  is the non-random discretization error. From the definition of  $Q(\delta)$  set observe that the event  $\{f(\hat{x}_n) - f(\hat{x}) > \delta\}$  is equivalent to  $\{\hat{x}_n \in Q(\delta)\}$ . Now, using the definition of  $\hat{x}_n$ , observe the equivalence

$$\{\hat{x}_n \in Q(\delta)\} = \bigcup_{x \in Q(\delta)} \bigcap_{y \in D} \{\hat{f}(x) \leq \hat{f}(y)\}.$$

Therefore, it straightforwardly follows that

$$\max_{x \in Q(\delta)} \mathbb{P} \left( \bigcap_{y \in D} \{\hat{f}(x) \leq \hat{f}(y)\} \right) \leq \mathbb{P}(f(\hat{x}_n) - f(x^*) > \varepsilon) \leq \sum_{x \in Q(\delta)} \mathbb{P} \left( \bigcap_{y \in D} \{\hat{f}(x) \leq \hat{f}(y)\} \right). \tag{5}$$

Now, since the set  $Q(\delta)$  is finite, Lemma 1.2.15 of Dembo and Zeitouni (2010) implies

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(f(\hat{x}_n) - f(x^*) > \varepsilon) &\leq \limsup_{n \rightarrow \infty} \frac{1}{n} \log \sum_{x \in Q(\delta)} \mathbb{P} \left( \bigcap_{y \in D} \{\hat{f}(x) \leq \hat{f}(y)\} \right) \\ &= \max_{x \in Q(\delta)} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left( \bigcap_{y \in D} \{\hat{f}(x) \leq \hat{f}(y)\} \right) \\ &\leq \max_{x \in Q(\delta)} \sum_{y \in D} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\hat{f}(x) \leq \hat{f}(y)), \end{aligned} \quad (6)$$

where the last inequality follows from the fact that sampling is independent at each of the design points. Next, for the lower bound, the monotonicity of the logarithm function and the sampling independence implies

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(f(\hat{x}_n) - f(x^*) > \varepsilon) &\geq \liminf_{n \rightarrow \infty} \frac{1}{n} \log \max_{x \in Q(\delta)} \mathbb{P} \left( \bigcap_{y \in D} \{\hat{f}(x) \leq \hat{f}(y)\} \right) \\ &\geq \max_{x \in Q(\delta)} \sum_{y \in D} \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\hat{f}(x) \leq \hat{f}(y)), \end{aligned} \quad (7)$$

where the final inequality follows from Lemmas 2 and 3. Finally, the theorem follows from an application of Lemma 1 to (6) and (7).  $\square$

Some comments are in order for this result. First, observe that the fact that we assume a finite grid  $D \subset \mathcal{X}$  implies that the LD rate function is well-defined. Furthermore, it can be anticipated that the rate function can be established even with a countable grid. Second, the set  $Q(\delta)$  is critical for establishing the limit. This is the set of design points in the grid  $D$  that are  $\varepsilon$  worse than the *optimal* design point  $x^* \in \mathcal{X}$ . Thus, the rate function  $J(\varepsilon)$  identifies the dominant point on the boundary of the set  $Q(\delta)$  that  $\hat{x}_n$  is most likely to diverge away from  $x^*$ . The form of  $J(\varepsilon)$  indicates that it is a composition of the closest point in  $Q(\delta)$  to each design point in  $D$  that is most likely to cause a mis-ordering.

#### 4 OPTIMAL LINEAR ALLOCATION RULE

In this section, we identify an (asymptotically) optimal linear allocation rule that optimizes the rate function in Theorem 1. Note that this optimization is *post hoc* in the sense that the rate function is identified for an arbitrary allocation. The optimal allocation maximizes the rate at which the log likelihood of misordering the empirical optimizer approaches zero. First, we show that the rate function obtained in Theorem 1 is strictly concave and thus has a unique maximizer. Let  $J(\alpha_x, \alpha_y, t) := \left( \alpha_y \Lambda \left( y, \frac{t}{\alpha_y} \right) + \alpha_x \Lambda \left( x, -\frac{t}{\alpha_x} \right) \right)$  for brevity.

**Lemma 4**  $I(\alpha_x, \alpha_y) = \sup_{t \in \mathbb{R}} \left( -\alpha_y \Lambda \left( y, \frac{t}{\alpha_y} \right) - \alpha_x \Lambda \left( x, -\frac{t}{\alpha_x} \right) \right)$ , is strictly concave  $\forall \{\alpha_x, \alpha_y\} \in [0, 1] \times [0, 1]$ .

*Proof.* Observe that proving  $I(\alpha_x, \alpha_y)$  is concave in  $\alpha_x$  and  $\alpha_y$  is equivalent to proving

$$\inf_{t \in \mathbb{R}} \left( \alpha_y \Lambda \left( y, \frac{t}{\alpha_y} \right) + \alpha_x \Lambda \left( x, -\frac{t}{\alpha_x} \right) \right) = \inf_{t \in \mathbb{R}} J(\alpha_x, \alpha_y, t)$$

is convex  $\forall \alpha_x, \alpha_y \in [0, 1] \times [0, 1]$ . First, we demonstrate that  $\alpha_y \Lambda\left(y, \frac{t}{\alpha_y}\right)$  is convex in  $t$  and  $\alpha_y$ . Using the definition of  $\Lambda(\cdot, t)$  we have

$$\frac{\partial^2}{\partial \alpha_y^2} \alpha_y \Lambda\left(y, \frac{t}{\alpha_y}\right) = \frac{t^2}{\alpha_y^3} \left\{ \frac{\mathbb{E}\left[\exp\left(t \frac{L(y, \xi)}{\alpha_y}\right) L(y, \xi)^2\right]}{\mathbb{E}\left[\exp\left(t \frac{L(y, \xi)}{\alpha_y}\right)\right]} - \frac{\mathbb{E}\left[\exp\left(t \frac{L(y, \xi)}{\alpha_y}\right) L(y, \xi)\right]^2}{\mathbb{E}\left[\exp\left(t \frac{L(y, \xi)}{\alpha_y}\right)\right]^2} \right\} > 0. \quad (8)$$

Observe that the expression on the right hand side of (8) is the variance of  $L(y, \xi)$ , with respect to the ‘twisted’ distribution,

$$\frac{\exp\left(t \frac{L(y, \xi)}{\alpha_y}\right) dF(\xi)}{\mathbb{E}\left[\exp\left(t \frac{L(y, \xi)}{\alpha_y}\right)\right]},$$

and the overall expression is strictly positive since  $t^2/\alpha_y^3 > 0$ . It follows that  $\alpha_y \Lambda\left(y, \frac{t}{\alpha_y}\right)$  is strictly convex in  $\alpha_y$ . Similarly, observe that  $\alpha_x \Lambda\left(x, -\frac{t}{\alpha_x}\right)$  is strictly convex in  $\alpha_x$ . Consequently, it is straightforward to see that the Hessian of  $J(\alpha_x, \alpha_y, t) \quad \forall t \in \mathbb{R}$  is positive definite. Therefore,  $J(\alpha_x, \alpha_y, t)$  is strictly convex  $\forall \alpha_x, \alpha_y \in [0, 1] \times [0, 1]$ . We also know that the cumulant generating function is convex in  $t$  in general, but it is strictly convex due to Assumption 1. Since the sum of two strictly convex function is strictly convex,  $J(\alpha_x, \alpha_y, t)$  is strictly convex  $\forall t \in \mathbb{R}$ .

Next, observe that for any  $\eta > 0$ , there exists a  $t \in \mathbb{R}$ , such that for a given  $\alpha_x, \alpha_y \in [0, 1] \times [0, 1]$ ,

$$J(\alpha_x, \alpha_y, t) \leq -I(\alpha_x, \alpha_y) + \eta.$$

We now follow the arguments in Sec. 3.2.5 of Boyd and Vandenberghe (2004). For any  $\beta \in [0, 1]$  and  $\alpha_x^j, \alpha_y^j \in [0, 1] \times [0, 1], j = 1, 2$ ,

$$\begin{aligned} -I(\beta \alpha_x^1 + (1 - \beta) \alpha_x^2, \beta \alpha_y^1 + (1 - \beta) \alpha_y^2) &= \inf_{t \in \mathbb{R}} J(\beta \alpha_x^1 + (1 - \beta) \alpha_x^2, \beta \alpha_y^1 + (1 - \beta) \alpha_y^2, t) \\ &\leq J(\beta \alpha_x^1 + (1 - \beta) \alpha_x^2, \beta \alpha_y^1 + (1 - \beta) \alpha_y^2, \beta t_1 + (1 - \beta) t_2) \\ &< \beta J(\alpha_x^1, \alpha_y^1, t_1) + (1 - \beta) J(\alpha_x^2, \alpha_y^2, t_2) \\ &\leq -\beta I(\alpha_x^1, \alpha_y^1) - (1 - \beta) I(\alpha_x^2, \alpha_y^2) + \eta, \end{aligned}$$

where penultimate inequality follows from Jensen’s inequality. Since  $\eta$  is arbitrary, it follows that  $I(\alpha_x, \alpha_y)$  is strictly concave.  $\square$

**Lemma 5**  $\min_{x \in Q(\delta)} \sum_{y \in D} I(\alpha_x, \alpha_y) \mathbb{I}_{\{f(y) < f(x)\}}$  is strictly concave in  $\{\alpha_1, \alpha_2, \dots, \alpha_d\} \in [0, 1]^d$ .

*Proof.* From Lemma 4 and the fact that the minimum of strictly concave functions preserves strict concavity, the proposition follows.  $\square$

**Theorem 2** The following constraint maximization problem is strictly concave,

$$\begin{aligned} \max_{\{\alpha_1, \alpha_2, \dots, \alpha_d\}} \quad & \min_{x \in Q(\delta)} \sum_{y \in D} I(\alpha_x, \alpha_y) \mathbb{I}_{\{f(y) < f(x)\}}, \quad (9) \\ \text{such that} \quad & \sum_{i=1}^d \alpha_i = 1, \quad \alpha_i \in [0, 1] \forall i \in \{1, 2, \dots, d\}. \end{aligned}$$

*Proof.* The proof immediately follows from Lemma 5.  $\square$

Hence, the optimal allocation strategy is the solution of (9). Next, we illustrate this optimization for specific cases.

**Example 1: Normal Distribution**

Assume that  $L(x, \xi) \sim \mathcal{N}(f(x), \sigma^2(x))$ , and observe that

$$\begin{aligned} I(\alpha_x, \alpha_y) &= \sup_{t \in \mathbb{R}} \left( -\alpha_y \Lambda \left( y, \frac{t}{\alpha_y} \right) - \alpha_x \Lambda \left( x, -\frac{t}{\alpha_x} \right) \right) \\ &= \frac{1}{2} (f(x) - f(y))^2 \left( \frac{\sigma^2(y)}{\alpha_y} + \frac{\sigma^2(x)}{\alpha_x} \right)^{-1}. \end{aligned} \tag{10}$$

Therefore, using Theorem 1, the LDP for a given  $\varepsilon > 0$  is

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(f(\hat{x}_n) - f(x^*) > \varepsilon) = - \min_{x \in Q(\delta)} \sum_{y \in D} \frac{1}{2} (f(x) - f(y))^2 \left( \frac{\sigma^2(y)}{\alpha_y} + \frac{\sigma^2(x)}{\alpha_x} \right)^{-1} \mathbb{I}_{\{f(x) > f(y)\}},$$

where  $Q(\delta)$  is as defined in Theorem 1.

**Example 2: Binomial Distribution**

Assume that  $L(x, \xi) \sim \text{Bin}(\frac{f(x)}{m}, m)$ , where  $m$  is the number of binomial trials. Observe that

$$\begin{aligned} I(\alpha_x, \alpha_y) &= \sup_{t \in \mathbb{R}} \left( -\alpha_y \Lambda \left( y, \frac{t}{\alpha_y} \right) - \alpha_x \Lambda \left( x, -\frac{t}{\alpha_x} \right) \right) \\ &= m \left( -\alpha_y \log \left( 1 - \frac{f(y)}{m} + \frac{f(y)}{m} e^{\frac{t^*}{\alpha_y}} \right) - \alpha_x \log \left( 1 - \frac{f(x)}{m} + \frac{f(x)}{m} e^{-\frac{t^*}{\alpha_x}} \right) \right), \end{aligned} \tag{11}$$

where  $t^* = \log \left( \frac{f(x)(m-f(y))}{f(y)(m-f(x))} \right) \left[ \frac{1}{\alpha_x} + \frac{1}{\alpha_y} \right]^{-1}$ . Therefore using Theorem 1, the LDP for a given  $\varepsilon > 0$  is

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\omega \in \Omega : f(\hat{x}_n) - f(x^*) > \varepsilon) &= - \min_{x \in Q(\delta)} \sum_{y \in D} m \left( -\alpha_y \log \left( 1 - \frac{f(y)}{m} + \frac{f(y)}{m} e^{\frac{t^*}{\alpha_y}} \right) \right. \\ &\quad \left. - \alpha_x \log \left( 1 - \frac{f(x)}{m} + \frac{f(x)}{m} e^{-\frac{t^*}{\alpha_x}} \right) \right) \mathbb{I}_{\{f(x) > f(y)\}}, \end{aligned}$$

where  $Q(\delta)$  is as defined in Theorem 1.

**5 SEQUENTIAL OPTIMIZATION**

The optimization problem (9) is solved by the master machine, and the sampling budget is assigned to the slave machines. Observe that (9) assumes that the master machine has complete knowledge of the true cumulant generating function of the (stochastic) loss function. In practice, of course, this is unknown and the master machine must rely on empirical estimates of the objective from the slave machines. An appropriate approach to solving the optimization problem would be to perform a sequential optimization as sample estimates accumulate. With a large, but finite budget the sequential optimum should be close to the optimizer of (9).

We demonstrate the computation in two different scenarios. First, we assume that the (rate) function  $J(\alpha_x, \alpha_y, t)$  can be analytically optimized over  $t$ . In this case, the objective in (9) is simpler to estimate and optimize. Note that these instances are direct analogues in the SAA context of the ranking and selection (R&S) problems studied in Glynn and Juneja (2004), Hunter and Pasupathy (2013), and Glynn and Juneja

(2018). To deal with these types of problems, we present Algorithm 1 below, that parallels Hunter and Pasupathy (2013) Algorithm 2.

Second, in many applications of SAA, the optimization of  $J(\cdot, \cdot, t)$  must be carried out numerically since closed forms are not available. These instances are far more complicated than the straightforward R&S analogues considered above, as the geometry of the loss function now plays a prominent role. Algorithm 2 below exploits an *expectation-maximization* (EM) type iterative structure to recursively compute the optimal allocation efficiently.

Of course, in either scenario, the rate functions and objectives must be estimated by Monte Carlo sampling at the slave machines. Let  $\hat{J}(\alpha_x, \alpha_y, t) := \alpha_y \hat{\Lambda}\left(y, \frac{t}{\alpha_y}\right) + \alpha_x \hat{\Lambda}\left(x, -\frac{t}{\alpha_x}\right)$ , where  $\hat{\Lambda}\left(x, \frac{t}{\alpha_x}\right) := ([n\alpha_x])^{-1} \sum_{i=1}^{[n\alpha_x]} \exp\left(\frac{t}{\alpha_x} L(x, \xi_i^x)\right)$  is the natural empirical estimator of the log moment generating function. We also define the set  $\hat{Q}(\delta) = \{x \in D : \hat{f}(x) > \hat{f}(x^*) + \varepsilon\}$ , for a given  $\varepsilon > 0$ . In the remainder of this section, we assume a fixed  $\varepsilon > 0$ . Let  $\hat{\alpha}_x(n)$  represent the estimated allocation at location  $x \in D$  with total sampling budget  $n$ . We define the *optimality gap* of the estimator as

$$OG(n) := \sum_{x \in D} |\hat{\alpha}_x^{(n)} - \alpha_x|,$$

where  $\alpha_x$  is the ‘true’ optimal linear allocation obtained by solving (9).

### 5.1 Optimization with Closed-forms

Consider situations where the ‘inner’ optimization  $\sup_{t \in \mathbb{R}} J(\alpha_x, \alpha_y, t)$   $x, y \in D$  can be completed in closed form analytically. For instance, in Examples 1 and 2 above closed forms were derived for cases where the loss functions at each of the design points are Gaussian and binomially distributed (respectively). Let  $\hat{I}(\cdot, \cdot)$  represent the Monte Carlo estimate of this closed form, which will require estimation of the mean and (possibly) the variance. Algorithm 1 proceeds iteratively by estimating the optimal allocation while accumulating more and more samples at each iteration. This method parallels Algorithm 2 in Hunter and Pasupathy (2013).

Table 1: Algorithm 1.

Step 0	Initialize pilot sample $n_x^{(0)} = N^0$ at each $x \in D$ .
For each $k \geq 0$ :	
Step 1	Generate $n_x^{(k)}$ i.i.d. samples at each $x \in D$ . Compute $\hat{I}(\alpha_x, \alpha_y) \forall x, y \in D, \hat{f}(x) \forall x \in D$ and $\hat{Q}(\delta)$ using all $\sum_{i=0}^k n_x^{(i)}$ samples.
Step 2	Compute $(\alpha_x^{(k+1)}, x \in D) = \arg \max_{\{\alpha_x, x \in D\}} \min_{\{x \in \hat{Q}(\delta)\}} \sum_{y \in D} \hat{I}(\alpha_x, \alpha_y) \mathbb{I}_{\{\hat{f}(x) > \hat{f}(y)\}}$ . Generate $T_j \forall j = \{1, 2, \dots, N\}$ , where $T_j$ has empirical distribution with probability $\alpha_x^{(k+1)}$ on support $x \in D$ . Set $n_x^{(k+1)} = \sum_{j=1}^N \mathbb{I}_{\{T_j=x\}} \forall x \in D$ .
Step 3	Repeat Steps 1 and 2 , until sampling budget exhausts.

Observe that the algorithm runs until the sampling budget is exhausted with no guarantees on convergence to the true optimal allocation rule. Consistency results from Hunter and Pasupathy (2013) and Glynn and Juneja (2004) imply that with a large, but finite budget the allocation obtained at the end of the procedure should closely match the optimal linear allocation rule. We now illustrate the algorithm by running through a couple of examples.

**Gaussian Loss:** For simplicity we assume that the variances  $\sigma^2(x) \forall x \in D$  are known, and the mean value at each of the design points is estimated using the natural estimator.

The choice of  $\delta$  determines the error tolerance, and affects the allocation budget. Figure 1(a) depicts a case where  $\delta$  is much smaller than the resolution of the grid. In this case, both the true allocation and the

estimated allocation place much of the sampling effort near the optimizer. On the other hand, when  $\delta$  is of the order of the grid resolution, Figure 1(b) demonstrates that both the true and the estimated allocations expend substantial sampling efforts near the “boundary” of the sets  $Q(\delta)$  and  $\hat{Q}(\delta)$  respectively. Figure 2 illustrates that in both the cases, optimality gap appears to converge, but with large variance.

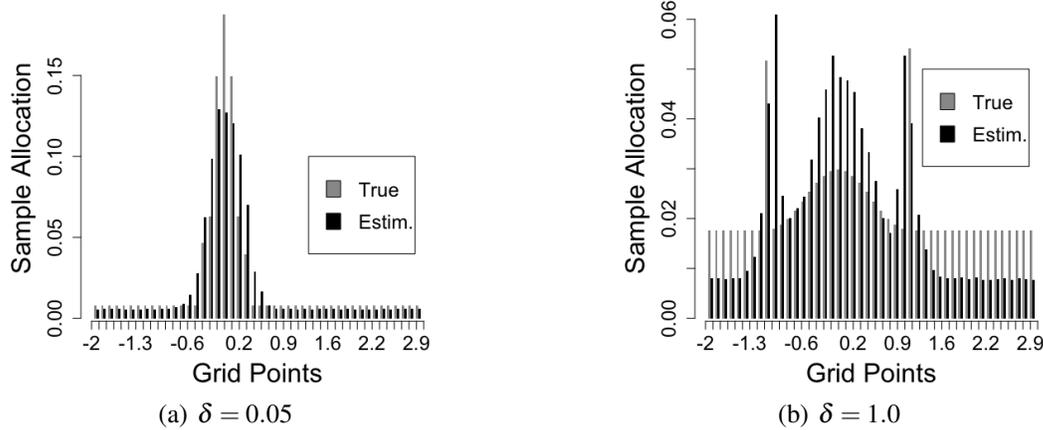


Figure 1: Approximate and true allocations with Gaussian distributed loss functions. Total sampling budget  $n = 4600$ ,  $|D| = 46$  and averaged over 50 sample paths.

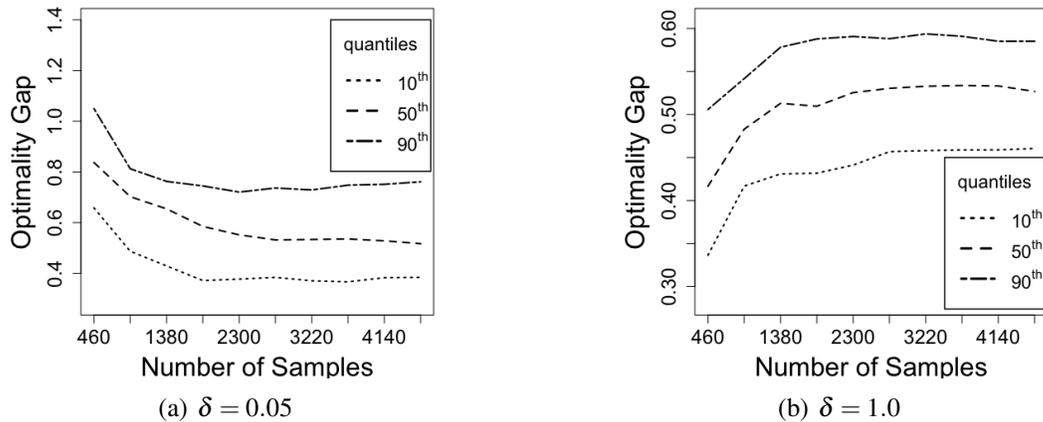


Figure 2: Sample quantile of Optimalty Gap for Gaussian distributed loss functions for 50 Sample paths. Total sampling budget  $n = 4600$ ,  $|D| = 46$ .

**Binomial Loss:** Next, in the case of the binomially distributed loss function, it suffices to compute the objective value  $\hat{f}(x)$  using the natural, plug-in estimator for (11). Our observations here parallel the Gaussian case.

### 5.2 Optimization without Closed-forms

It is rare to place explicit distributional assumptions on the loss function at different design points in the grid, and typical stochastic programming models assume regularity conditions on the loss function  $L(\cdot, \cdot)$  and stochasticity conditions on  $\xi$ . The distributional conditions are then consequences of these two ingredients. In general, then, it is highly unlikely that there exists a closed form for the optimization  $\sup_{t \in \mathbb{R}} J(\cdot, \cdot, t)$ , and the optimization must be carried out numerically on a Monte Carlo estimate  $\hat{J}(\cdot, \cdot, t)$ .

As noted before, Algorithm 1 has no guarantees on convergence within a fixed number of iterations, since it is only running till the sampling budget is exhausted. On the other hand, there are many applications of SAA where it is useful to run the algorithm till convergence. For instance, in data-driven applications, it may be possible to obtain  $n$  samples repeatedly from a simulator, or by bootstrap sampling of a given dataset. To handle such situations we propose a second iterative algorithm that is expectation-maximization (EM) like, and proceeds in two iterative steps. In step one, for a fixed linear allocation rule  $(\alpha_x, x \in D)$ , we compute  $\hat{J}(\alpha_x, \alpha_y, t)$  for every  $x, y \in D$  and identify the optimal  $t(x, y)$ . In step two, we compute an allocation using the objective in 9 albeit with  $I(\alpha_x, \alpha_y, t(x, y))$  from step one. These two steps are iterated till there is no improvement in the allocation in step two. The following display summarizes the algorithm.

Table 2: Algorithm 2.

Step 0	Fix $(\alpha_x^{(0)}, x \in D)$ , where $\alpha_x^{(0)} \in [0, 1]$ and $\sum_x \alpha_x^{(0)} = 1$ .
For each $k \geq 0$ :	
Step 1	Generate $T_j \forall j = \{1, 2, \dots, n\}$ , where $T_j$ has empirical distribution with probability $\alpha_x^{(k)}$ on support $x \in D$ . Set $n_x^{(k)} = \sum_{j=1}^n \mathbb{I}_{\{T_j=x\}} \forall x \in D$ . Generate $n_x^{(k)}$ i.i.d. random samples at each $x \in D$ . Compute $t^{(k)}(x, y) = \arg \sup_{t \in \mathbb{R}} \hat{J}(\alpha_x^{(k)}, \alpha_y^{(k)}, t)$ for all $x, y \in D$ , and Compute $\hat{f}(x)$ for all $x \in D$ and $\hat{Q}(\delta)$ using all $\sum_{i=0}^k n_x^{(i)}$ samples.
Step 2	Compute $(\alpha_x^{(k+1)}, x \in D) \in \arg \max_{\{\alpha_x, x \in D\}} \min_{x \in \hat{Q}(\delta)} \sum_{y \in D} \hat{I}(\alpha_x, \alpha_y, t^{(k)}(x, y)) \mathbb{I}_{\{\hat{f}(x) > \hat{f}(y)\}}$ .
Step 3	Repeat Step 1 and 2, until $(\alpha_x^{(k)}, x \in D)$ converges.

Note that we assume that the sampling budget  $n$  is fixed, and allow the algorithm to produce  $n$  samples on each iteration. It is possible to couple the iterative scheme to the sampling budget by fixing the number of iterations *a priori* to  $K = \lfloor n * \gamma \rfloor$ , where  $\gamma \in [0, 1]$  is fixed. At the  $k$ th iteration,  $\sum_{i=0}^k \alpha_x^{(i)} K$  samples is generated at sample point  $x \in D$ . In this case, the algorithm terminates once  $K$  iterations have been completed.

**Mean-squared Error** Consider a squared error loss, widely used in empirical risk minimization of machine learning models, where  $L(x, \xi) := (x - \xi)^2$ . For simplicity, we assume that the design space is one dimensional and that  $\xi$  is Gaussian.

Observe that the estimated allocation is quite close to the true allocation, even in this case (see Figure 3). Significant budget allocations are once again made in the vicinity of the boundary of  $Q(\delta)$ . Figure 4 on the other hand, amply demonstrates that while the estimators are consistent with a large budget “on average” as demonstrated by the 50th percentile lines, there is significant variance in the optimality gap even at large sample values as shown by the spread between the 90th and 10th percentiles.

## 6 CONCLUSIONS AND FUTURE DIRECTIONS

We study the problem of optimally allocating a sampling budget in order to compute a sample average approximation (SAA) of the solution of a single-stage stochastic program. Under a fixed finite discretization of the design space (or ‘grid’), we first establish a large deviations principle satisfied by the regret, defined as  $f(\hat{x}_n) - f(x^*)$ , where  $f(\cdot)$  is the true objective,  $\hat{x}_n$  is the SAA estimate of the optimizer and  $x^*$  the true optimizer. Next, we identify a constraint maximization problem, whose solution identifies an optimal linear allocation rule that maximizes the decay rate of the likelihood of identifying an incorrect optimal design point, in the limit of a large sampling budget. Finally, we designed two different algorithms to sequentially implement this optimization.

The developments in this paper lead to multiple important and open problems, relevant to both simulation optimization and machine learning more broadly. First, our current treatment of the regret effectively assumes

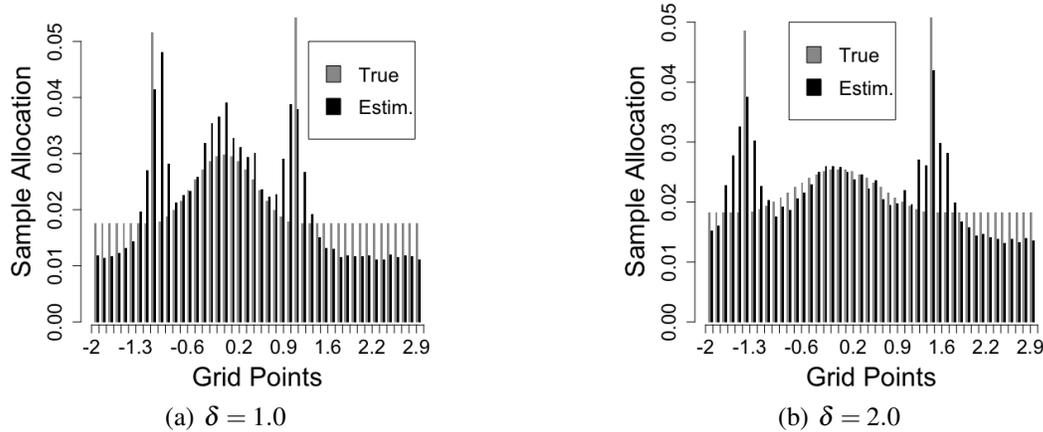


Figure 3: Approximate and true allocations for squared loss with Gaussian random samples. Sample at each iteration  $n = 460$ ,  $|D| = 46$  and averaged over 30 sample paths.

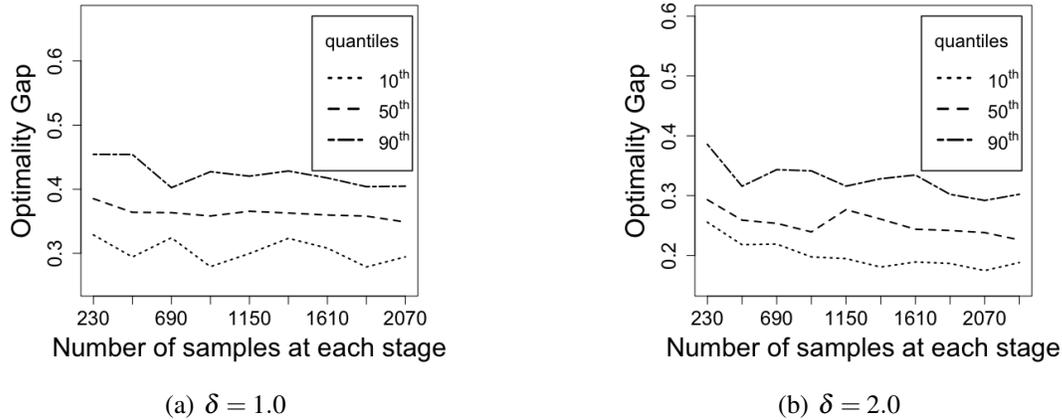


Figure 4: Sample quantile of Optimality Gap for squared loss functions with Gaussian random samples for 30 Sample paths.  $|D| = 46$ .

that the grid is fixed. An important question is how the grid size affects the large deviations rate function. In particular, it can be easily seen that when the grid has the cardinality of the continuum, the rate function does not exist. On the other hand, there is definite benefit in scaling the grid size with the sampling budget. How should this be done to obtain a large deviations principle in the limit?

Second, while results in Glynn and Juneja (2004) and Hunter and Pasupathy (2013) can be straightforwardly adapted to establish consistency of the sequentially estimated allocation rule, the efficiency of the estimator is unknown. In particular, we conjecture that the rate function estimators used are highly inefficient. This follows from the fact that we use the canonical estimator for the cumulant generating function, and it is conjectured that the latter estimators are heavy-tailed (Glynn and Juneja 2011). On the other hand, we are really only interested in the accuracy of the estimated objective in the vicinity of the true optimizer, and not the global accuracy. Closer to the optimizer, and in the large sampling budget limit, we conjecture that it is possible to use fewer moments to accurately estimate the rate function, leading to substantial improvements in efficiency.

Third, note that the algorithms designed here are not dimension free, and we conjecture that even with strictly convex objective functions  $f$  they will not scale well. We postulate that it is possible to combine (9)

with multiple stochastic gradient descent (SGD) crawlers starting at each of the grid-points, and letting these iterate a fixed number of times, to make the allocation optimization algorithm dimension-free.

Fourth, ‘gridding’ the design space has significant algorithmic advantages since it allows our present algorithms to be implemented in a ‘divide and conquer’ manner, whereby slave machines sample and estimate moments of the loss function at each design point, and communicate this to the master machine for optimization. On the other hand, estimates at a given design point must be close to those at other design points in its neighborhood. With a fine grid, the amount of communication overhead required to implement a fully parallel computation is likely significant. We speculate that it should be possible to design a completely decentralized optimization scheme to determine the optimal allocation just by performing local message passing. Such a scheme would yield substantial reductions in communication overhead.

## REFERENCES

- Boyd, S., and L. Vandenberghe. 2004. *Convex Optimization*. New York: Cambridge University Press.
- Chen, C.-H., J. Lin, E. Yücesan, and S. E. Chick. 2000. “Simulation Budget Allocation for further enhancing the Efficiency of Ordinal Optimization”. *Discrete Event Dynamic Systems* 10(3):251–270.
- Dembo, A., and O. Zeitouni. 2010. *Large Deviations Techniques and Applications, volume 38 of Stochastic Modelling and Applied Probability*. 2nd ed. Berlin: Springer-Verlag.
- Glynn, P., and S. Juneja. 2004. “A Large Deviations Perspective on Ordinal Optimization”. In *Proceedings of the 2004 Winter Simulation Conference*, edited by R. G. Ingalls et al., 577–585. Piscataway, New Jersey: IEEE.
- Glynn, P., and S. Juneja. 2011. “Ordinal Optimization: A Nonparametric framework”. In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain et al., 4062–4069. Piscataway, New Jersey: IEEE.
- Glynn, P., and S. Juneja. 2018. “Selecting the Best System, Large Deviations, and Multi-armed Bandits”. *arXiv preprint arXiv: 1507.04564v2*.
- Hunter, S. R., and R. Pasupathy. 2013. “Optimal Sampling Laws for Stochastically Constrained Simulation Optimization on Finite Sets”. *INFORMS Journal on Computing* 25(3):527–542.
- Shapiro, A., D. Dentcheva, and A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. 2nd ed. Philadelphia, Pennsylvania: SIAM.

## AUTHOR BIOGRAPHIES

**PRATEEK JAISWAL** is a Ph.D. student in School of Industrial Engineering at Purdue University. His research interests are in stochastic optimization. His e-mail address is [jaiswalp@purdue.edu](mailto:jaiswalp@purdue.edu).

**RAGHU PASUPATHY** is an associate professor in the Department of Statistics at Purdue University. His research interests lie broadly in Monte Carlo methods with a specific focus on simulation optimization. He is a member of INFORMS, IIE, and ASA, and serves as an associate editor for Operations Research and INFORMS Journal on Computing. His email address is [pasupath@purdue.edu](mailto:pasupath@purdue.edu).

**HARSHA HONNAPPA** is an assistant professor in the School of Industrial Engineering at Purdue University. His research interests are in queueing theory, applied probability, game theory, and statistical modeling. He is also the recipient of the 2016 Lajos Takacs award for outstanding Ph.D. thesis on Queueing Theory and its Applications. His email address is [honnappa@purdue.edu](mailto:honnappa@purdue.edu).