

Accepted in *Operations Research*

Robust Analysis in Stochastic Simulation: Computation and Performance Guarantees

Soumyadip Ghosh

IBM Research AI, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, ghosh.s@us.ibm.com

Henry Lam

Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027,
henry.lam@columbia.edu

Any performance analysis based on stochastic simulation is subject to the errors inherent in misspecifying the modeling assumptions, particularly the input distributions. In situations with little support from data, we investigate the use of worst-case analysis to analyze these errors, by representing the partial, nonparametric knowledge of the input models via optimization constraints. We study the performance and robustness guarantees of this approach. We design and analyze a numerical scheme for solving a general class of simulation objectives and uncertainty specifications. The key steps involve a randomized discretization of the probability spaces, a simulable unbiased gradient estimator using a nonparametric analog of the likelihood ratio method, and a Frank-Wolfe (FW) variant of the stochastic approximation (SA) method (which we call FWSA) run on the space of input probability distributions. A convergence analysis for FWSA on non-convex problems is provided. We test the performance of our approach via several numerical examples.

1. Introduction

Simulation-based performance analysis of stochastic models, or stochastic simulation, is built on input model assumptions that to some extent deviate from the truth. Consequently, a performance analysis subject to these input errors may lead to poor prediction and suboptimal decision-making. To address this important problem, a typical framework in the stochastic simulation literature focuses on output variability measures or confidence bounds that account for the input uncertainty when input data are available. Established statistical techniques such as the bootstrap (e.g., Barton and Schruben (1993), Barton et al. (2013)), goodness-of-fit tests (e.g., Banks et al. (2009)), Bayesian

inference and model selection (e.g., Chick (2001), Zouaoui and Wilson (2004)) and the delta method (e.g., Cheng and Holland (1998, 2004)) have been proposed and have proven effective in many situations.

In this paper, we take a different approach for situations with insufficient data, or when the modeler wants to assess risk beyond what the data or the model indicates. Such situations can arise when the system, service target or operational policy in study is at a testing stage without much prior experience. To find reliable output estimates in these settings, we investigate a worst-case approach with respect to the input models. In this framework, the modeler represents the partial and nonparametric beliefs about the input models as constraints, and computes tight worst-case bounds among all models that satisfy them. More precisely, let $Z(P^1, \dots, P^m)$ be a performance measure that depends on m input models, each generated from a probability distribution P^i . The formulation for computing the worst-case bounds are

$$\min_{P^i \in \mathcal{U}^i, i=1, \dots, m} Z(P^1, \dots, P^m) \quad \text{and} \quad \max_{P^i \in \mathcal{U}^i, i=1, \dots, m} Z(P^1, \dots, P^m) \quad (1)$$

The set \mathcal{U}^i encodes the collection of all possible P^i from the knowledge of the modeler. The decision variables in the optimizations in (1) are the unknown models $P^i, i = 1, \dots, m$.

The primary motivation for using (1) is the robustness against model misspecification, where a proper construction of the set \mathcal{U}^i avoids making specific assumptions beyond the modeler's knowledge. The following three examples motivate and explain further.

EXAMPLE 1 (ROBUST BOUNDS UNDER EXPERT OPINION). When little information is available for an input model, a common practice in stochastic simulation is to summarize its range (say $[a, b]$) and mean (or mode) as a triangular distribution, where the base of the triangle denotes the range and the position of the peak is calibrated from the mean. This specific distribution only crudely describes the knowledge of the modeler and may deviate from the true distribution, even if a, b, μ are correctly specified. Instead, using

$$\mathcal{U}^i = \{P^i : E_{P^i}[X^i] = \mu, \text{ supp } P^i = [a, b]\} \quad (2)$$

in formulation (1), where X^i is the random variate, $E_{P^i}[\cdot]$ is the expectation under P^i , and $\text{supp } P^i$ is the support of P^i , will give a valid interval that covers the true performance measure whenever a, b, μ are correctly specified. Moreover, when these parameters are not fully known but instead specified within a range, (2) can be relaxed to

$$\mathcal{U}^i = \{P^i : \underline{\mu} \leq E_{P^i}[X^i] \leq \bar{\mu}, \text{ supp } X^i = [\underline{a}, \bar{b}]\}$$

where $[\underline{\mu}, \bar{\mu}]$ denotes the range of the mean and \underline{a}, \bar{b} denote the lower estimate of the lower support end and upper estimate of the upper support end respectively. The resulting bound will cover the truth as long as these ranges are supplied correctly. \square

EXAMPLE 2 (DEPENDENCY MODELING). In constructing dependent input models, common approaches in the simulation literature fit the marginal description and the correlation of a multivariate model to a specified family. Examples include Gaussian copula (e.g., Lurie and Goldberg (1998), Charnouf and L'Ecuyer (2009); also known as normal-to-anything (NORTA), e.g. Cario and Nelson (1997)) and chessboard distribution (Ghosh and Henderson (2002)) that uses a domain discretization. These distributions are correctly constructed up to their marginal description and correlation, provided that these information are correctly specified. However, dependency structure beyond correlation can imply errors on these approaches (e.g., Lam (2017)), and formulation (1) can be used to get bounds that address such dependency. For example, suppose P^i is a bivariate input model with marginal distributions $P^{i,1}, P^{i,2}$, marginal means $\mu^{i,1}, \mu^{i,2}$ and covariance ρ^i . We can set

$$\mathcal{U}^i = \{P^i : P_{P^{i,1}}(X^{i,1} \leq q_j^{i,1}) = \nu_j^1, j = 1, \dots, l_1, P_{P^{i,2}}(X^{i,2} \leq q_j^{i,2}) = \nu_j^2, j = 1, \dots, l_2, E[X^{i,1}X^{i,2}] = \rho^i + \mu^{i,1}\mu^{i,2}\}$$

where $(X^{i,1}, X^{i,2})$ denote the random vector under P^i , and $q_j^{i,1}, q_j^{i,2}, \nu_j^{i,1}, \nu_j^{i,2}$ are pre-specified quantiles and probabilities of the respective marginal distributions. Unlike previous approaches, (1) outputs correct bounds on the truth given correctly specified marginal quantiles and correlation, regardless of the dependency structure. \square

EXAMPLE 3 (MODEL RISK). Model risk refers broadly to the uncertainty in analysis arising from the adopted model not being fully accurate. This inaccuracy occurs as the adopted model (often known as the baseline model), typically obtained from the best statistical fit or expert opinion, deviates from the truth due to the real-world non-stationarity and the lack of full modeling knowledge or capability. To assess model risk, a recently surging literature studies the use of statistical distance as a measurement of model discrepancy (e.g., Glasserman and Xu (2014), Lam (2016b)). Given the baseline model P_b^i , the idea is to represent the uncertainty in terms of the distance away from the baseline via a neighborhood ball

$$\mathcal{U}^i = \{P^i : d(P^i, P_b^i) \leq \eta^i\} \quad (3)$$

where d is a distance defined on the nonparametric space of distributions (i.e., without restricting to any parametric families). The bounds drawn from formulation (1) assess the effect of model risk due to the input models, tuned by the ball size parameter η^i that denotes the uncertainty level. Besides risk assessment, this approach can also be used to obtain consistent confidence bounds for the true performance measure, when P_b^i is taken as the empirical distribution and η and d are chosen suitably (discussed further in Section 3). \square

Our worst-case approach is inspired from the literature of robust optimization (Ben-Tal et al. (2009), Bertsimas et al. (2011)), which considers decision-making under uncertainty and advocates optimizing decisions over worst-case scenarios. In particular, when the uncertainty lies in the probability distributions that govern a stochastic problem, the decision is made to optimize under the worst-case distributions, a class of problems known as distributionally robust optimization (e.g. Delage and Ye (2010), Lim et al. (2006)). Such an approach has also appeared in so-called robust simulation or robust Monte Carlo in the simulation literature (Hu et al. (2012), Glasserman and Xu (2014)). However, the methodologies presented in the above literature focus on structured problems where the objective function is tractable, such as linear or linearly decomposable. In contrast, $Z(\cdot)$ for most problems in stochastic simulation is nonlinear and unstructured, obstructing the direct

adaptation of the existing methods. In view of this, our main objective is to design an efficient simulation-based method to compute the worst-case bounds for formulation (1) that can be applied to broad classes of simulation models and input uncertainty representations.

1.1. Our Contributions

We study a simulation-based iterative procedure for the worst-case optimizations (1), based on a modified version of the celebrated stochastic approximation (SA) method (e.g. Kushner and Yin (2003)). Because of the iterative nature, it is difficult to directly operate on the space of continuous distributions except in very special cases. Thus, our first contribution (Section 3) is to provide a randomized discretization scheme that can provably approximate the continuous counterpart. This allows one to focus on discrete distributions on fixed support points as the decision variable to feed into our SA algorithm.

We develop the SA method in several aspects. In Section 4, we construct an unbiased gradient estimator for Z based on the idea of the Gateaux derivative for functionals of probability distributions (Serfling (2009)), which is used to obtain the direction in each subsequent SA iterate. The need for such a construction is motivated by the difficulty in naïve implementation of standard gradient estimators: An arbitrary perturbation of a probability distribution, which is the decision variable in the optimization, may shoot outside the probability simplex and results in a gradient that does not bear any probabilistic meaning and subsequently does not support simulation-based estimation. Our approach effectively restricts the direction of perturbation to points within the probability simplex, leading to a simulable gradient estimator. We justify our approach as a non-parametric version of the classical likelihood ratio method (or the score function method) (Glynn (1990), Reiman and Weiss (1989), Rubinstein (1986)).

Next, in Sections 5 and 6, we design and analyze our SA scheme under the uncertainty constraints. We choose to use a stochastic counterpart of the so-called Frank-Wolfe (FW) method (Frank and Wolfe (1956)), known synonymously as the conditional gradient method in deterministic nonlinear programming. For convenience we call our scheme FWSA. Note that a standard SA

iteration follows the estimated gradient up to a pre-specified step size to find the next candidate iterate. When the formulation includes constraints, the common approach in the SA literature projects the candidate solution onto the feasible region in order to define the next iterate (e.g. Kushner and Yin (2003)). Instead, our method looks in advance for a feasible direction along which the next iterate is guaranteed to lie in the (convex) feasible region. In order to find this feasible direction, an optimization subproblem with a linear objective function is solved in each iteration. We base our choice of using FWSA on its computational benefit in solving these subproblems, as their linear objectives allow efficient solution scheme for high-dimensional decision variables for many choices of the set \mathcal{U}^i .

We characterize the convergence rate of FWSA in terms of the step size and the number of simulation replications used to estimate the gradient at each iteration. The form of our convergence bounds suggests prescriptions for the step-size and sample-size sequences that are efficient with respect to the cumulative number of sample paths simulated to generate all the gradients until the current iterate. The literature on the stochastic FW methods for non-convex problems is small. Kushner (1974) proves almost sure convergence under assumptions that can prescribe algorithmic specifications only for one-dimensional settings. During the review process of this paper, two other convergence rate studies Reddi et al. (2016) and Lafond et al. (2016) have appeared. Both of them assume the so-called G -Lipschitz condition on the gradient estimator that does not apply to our setting. Consequently, our obtained convergence rates are generally inferior to their results. Nonetheless, we will point out how our rates almost match theirs under stronger assumptions on the behavior of the iterates that we will discuss.

Finally, in Section 7 we provide numerical validation of our approach using two sets of experiments, one testing the performance of our proposed randomized discretization strategy, and one on the convergence of FWSA.

1.2. Literature Review

We briefly survey three lines of related work. First, our paper is related to the literature on input model uncertainty. In the parametric regime, studies have focused on the construction of confidence

intervals or variance decompositions to account for both parameter and stochastic uncertainty using data, via for instance the delta method (Cheng and Holland (1998, 2004)), the bootstrap (Barton et al. (2013), Cheng and Holland (1997)), Bayesian approaches (Zouaoui and Wilson (2003), Xie et al. (2014), Saltelli et al. (2010, 2008)), and metamodel-assisted analysis (Xie et al. (2014, 2015)). Model selection beyond a single parametric model can be handled through goodness-of-fit or Bayesian model selection and averaging (Chick (2001), Zouaoui and Wilson (2004)). Fully nonparametric approaches using the bootstrap have also been investigated (Barton and Schruben (1993, 2001), Song and Nelson (2015)).

Second, formulation (1) relates to the literature on robust stochastic control (Petersen et al. (2000), Iyengar (2005), Nilim and El Ghaoui (2005), Xu and Mannor (2012)) and distributionally robust optimization (Delage and Ye (2010), Goh and Sim (2010), Ben-Tal et al. (2013), Wiesemann et al. (2014)), where the focus is to make decision rules under stochastic environments that are robust against the ambiguity of the underlying probability distributions. This is usually cast in the form of a minimax problem where the inner maximization is over the space of distributions. This idea has spanned across multiple areas like economics (Hansen and Sargent (2001, 2008)), finance (Glasserman and Xu (2013), Lim et al. (2011)), queueing (Bertsimas and Natarajan (2007), Jain et al. (2010)), dynamic pricing (Lim and Shanthikumar (2007)), inventory management (Xin and Goldberg (2015)), physical sciences (Dupuis et al. (2016)), and more recently machine learning (Shafieezadeh-Abadeh et al. (2015), Blanchet et al. (2016)). In the simulation context, Hu et al. (2012) compared different global warming policies using Gaussian models with uncertain mean and covariance information. Glasserman and Xu (2014), Glasserman and Yang (2016) studied approaches based on sample average approximation for solving distance-based constrained optimizations to quantify model risk in finance. Lam (2016b, 2017) investigated infinitesimal approximations for related optimizations to quantify model errors arising from sequences of uncertain input variates. Bandi and Bertsimas (2012) studied the view of deterministic robust optimization to compute various stochastic quantities. Simulation optimization under input uncertainty

has also been studied via the robust optimization framework (Fan et al. (2013), Ryzhov et al. (2012)), and the closely related approach using risk measures (Qian et al. (2015), Zhou and Xie (2015)). Lastly, optimizations over probability distributions have also arisen as generalized moment problems, applied to decision analysis (Smith (1995, 1993), Bertsimas and Popescu (2005)) and stochastic programming (Birge and Wets (1987)).

Our algorithm relates to the literature on the FW method (Frank and Wolfe (1956)) and constrained SA. The former is a nonlinear programming technique initially proposed for convex optimization, based on sequential linearization of the objective function using the gradient at the solution iterate. The classical work of Canon and Cullum (1968), Dunn (1979) and Dunn (1980) analyzed convergence properties of FW for deterministic convex programs. More recently, Jaggi (2013), Freund and Grigas (2014) and Hazan and Luo (2016) carried out finite-time analysis for the FW method motivated by machine learning applications. For stochastic FW on non-convex problems (viewed as a class of constrained SA), Kushner (1974) focused on almost sure convergence based on a set of assumptions about the probabilistic behavior of the iterations, which were then used to tune the algorithm for one-dimensional problems. While this paper was under review, Reddi et al. (2016) provided a complexity analysis in terms of the sample size in estimating gradients and the number of calls of the linear optimization routine. Lafond et al. (2016) studied the performance in terms of regret in an online setting. Both Reddi et al. (2016) and Lafond et al. (2016) relied on the G -Lipschitz condition that our gradient estimator violated. Other types of constrained SA schemes include the Lagrangian method (Buche and Kushner (2002)) and mirror descent SA (Nemirovski et al. (2009)). Lastly, general convergence results for SA can be found in Fu (1994), Kushner and Yin (2003) and Pasupathy and Kim (2011).

2. Formulation and Assumptions

We focus on $Z(P^1, \dots, P^m)$ that is a finite horizon performance measure generated from i.i.d. replications from the independent input models P^1, \dots, P^m . Let $\mathbf{X}^i = (X_t^i)_{t=1, \dots, T^i}$ be T^i i.i.d.

random variables on the space $\mathcal{X}^i \subset \mathbb{R}^{u^i}$, each generated under P^i . The performance measure can be written as

$$Z(P^1, \dots, P^m) = E_{P^1, \dots, P^m}[h(\mathbf{X}^1, \dots, \mathbf{X}^m)] = \int \dots \int h(\mathbf{x}^1, \dots, \mathbf{x}^m) \prod_{t=1}^{T^1} dP(x_t^1) \dots \prod_{t=1}^{T^m} dP(x_t^m) \quad (4)$$

where $h(\cdot) : \prod_{i=1}^m (\mathcal{X}^i)^{T^i} \rightarrow \mathbb{R}$ is a cost function, and $E_{P^1, \dots, P^m}[\cdot]$ denotes the expectation associated with the generation of the i.i.d. replications. We assume that $h(\cdot)$ can be evaluated by the computer given the inputs. In other words, the performance measure (4) can be approximated by running simulation.

(4) is the stylized representation for transient performance measures in discrete-event simulation. For example, \mathbf{X}^1 and \mathbf{X}^2 can be the sequences of interarrival and service times in a queue, and P^1 and P^2 are the interarrival time and service time distributions. When $h(\mathbf{X}^1, \mathbf{X}^2)$ is the indicator function of the waiting time exceeding a threshold, (4) will denote the corresponding threshold exceedance probability.

Next we discuss the constraints in (1). Following the terminology in robust optimization, we call \mathcal{U}^i the *uncertainty set* for the i -th input model. Motivated by the examples in the Introduction, we focus on two types of convex uncertainty sets:

1. *Moment and support constraints:* We consider

$$\mathcal{U}^i = \{P^i : E_{P^i}[f_l^i(X^i)] \leq \mu_l^i, l = 1, \dots, s^i, \text{ supp } P^i = A^i\} \quad (5)$$

where X^i is a generic random variable under distribution P^i , $f_l^i : \mathcal{X}^i \rightarrow \mathbb{R}$, and $A^i \subset \mathcal{X}^i$. For instance, when $\mathcal{X}^i = \mathbb{R}$, $f_l^i(x)$ being x or x^2 denotes the first two moments. When $\mathcal{X}^i = \mathbb{R}^2$, $f_l^i(x_1, x_2) = x_1 x_2$ denotes the cross-moment. Equalities can also be represented via (5) by including $E_{P^i}[-f_l^i(X^i)] \leq -\mu_l^i$. Thus the uncertainty set (5) covers Examples 1 and 2 in the Introduction.

Furthermore, the neighborhood measured by certain types of statistical distance (Example 3) can also be cast as (5). For instance, suppose d is induced by the sup-norm on the distribution function on \mathbb{R} . Suppose P^i is a continuous distribution and the baseline distribution P_b^i is discrete with support points $y_j, j = 1, \dots, n^i$. The constraint

$$\sup_{x \in \mathbb{R}} |F^i(x) - F_b^i(x)| \leq \eta^i \quad (6)$$

where F^i and F_b^i denote the distribution functions for P^i and P_b^i respectively, can be reformulated as

$$F_b^i(y_j+) - \eta^i \leq F^i(y_j) \leq F_b^i(y_j-) + \eta^i, \quad j = 1, \dots, n^i$$

where $F_b^i(y_j-)$ and $F_b^i(y_j+)$ denote the left and right limits of F_b^i at y_j , by using the monotonicity of distribution functions. Thus

$$\mathcal{U}^i = \{P^i : F_b^i(y_j+) - \eta^i \leq E^i[I(X^i \leq y_j)] \leq F_b^i(y_j-) + \eta^i, \quad j = 1, \dots, n^i, \text{ supp } P^i = \mathbb{R}\}$$

where $I(\cdot)$ denotes the indicator function, falls into the form of (5). Bertsimas et al. (2014) considers this reformulation for constructing uncertainty sets for stochastic optimization problems, and suggests to select η^i as the quantile of the Kolmogorov-Smirnov statistic if F_b^i is the empirical distribution function constructed from continuous i.i.d. data.

2. *Neighborhood of a baseline model measured by ϕ -divergence:* Consider

$$\mathcal{U}^i = \{P^i : d_\phi(P^i, P_b^i) \leq \eta^i\} \quad (7)$$

where $d_\phi(P^i, P_b^i)$ denotes the ϕ -divergence from a baseline distribution P_b^i given by

$$d_\phi(P^i, P_b^i) = \int \phi \left(\frac{dP^i}{dP_b^i} \right) dP_b^i$$

which is finite only when P^i is absolutely continuous with respect to P_b^i . The function ϕ is a convex function satisfying $\phi(1) = 0$. This family covers many widely used distances. Common examples are $\phi(x) = x \log x - x + 1$ giving the KL divergence, $\phi(x) = (x - 1)^2$ giving the (modified) χ^2 -distance, and $\phi(x) = (1 - \theta + \theta x - x^\theta)/(\theta(1 - \theta))$, $\theta \neq 0, 1$ giving the Cressie-Read divergence. Details of ϕ -divergence can be found in, e.g., Pardo (2005), Ben-Tal et al. (2013), Bayraksan and Love (2015).

As precursed in the Introduction, in the context of simulation analysis where (P^1, \dots, P^m) are the input models, $Z(\cdot)$ in (4) is in general a complex nonlinear function. This raises challenges in solving (1) beyond the literature of robust control and optimization that considers typically more tractable objectives. Indeed, if $Z(\cdot)$ is a linear function in P^i 's, then optimizing over the two types of uncertainty sets above can both be cast as specialized classes of convex programs that

can be efficiently solved. But linear $Z(\cdot)$ is too restrictive to describe the input-output relation in simulation. To handle a broader class of $Z(\cdot)$ and to address its simulation-based nature, we propose to use a stochastic iterative method. The next sections will discuss our methodology in relation to the performance guarantees provided by (1).

3. Performance Guarantees and Discretization Strategy

This section describes the guarantees provided by our framework. Section 3.1 first presents the motivation and justification of a discretization scheme for continuous input distributions. Section 3.2 then discusses the statistical implications in more details.

3.1. Randomized Discretization

Suppose there is a “ground true” distribution P_0^i for each input model. Let Z_* and Z^* be the minimum and maximum values of the worst-case optimizations (1). Let Z_0 be the true performance measure, i.e. $Z_0 = Z(P_0^1, \dots, P_0^m)$. The following highlights an immediate implication of using (1):

PROPOSITION 1. *If $P_0^i \in \mathcal{U}^i$ for all i , then $Z_* \leq Z_0 \leq Z^*$.*

In other words, the bounds from the worst-case optimizations form an interval that covers the true performance measure if the uncertainty sets contain the true distributions.

We discuss a discretization strategy for the worst-case optimizations for continuous input distributions. We will show that, by replacing the continuous distribution with a discrete distribution on support points that are initially sampled from some suitably chosen distribution, we can recover the guarantee in Proposition 1 up to a small error. The motivation for using discretization comes from the challenges in handling decision variables in the form of continuous distributions when running our iterative optimization scheme proposed later.

We focus on the two uncertainty sets (5) and (7). The following states our guarantee:

THEOREM 1. *Consider $Z(P^1, \dots, P^m)$ in (4). Assume h is bounded a.s.. Let $n^i, i = 1, \dots, m$ and n be positive integers such that $n^i = nw^i$ for some fixed $w^i > 0$, for all i . For each input model i , we sample n^i i.i.d. observations $\{y_1^i, \dots, y_{n^i}^i\}$ from a distribution Q^i such that the true distribution P_0^i*

is absolutely continuous with respect to Q^i , with $L^i = dP_0^i/dQ^i$ satisfying $\|L^i\|_\infty < \infty$, where $\|L^i\|_\infty$ denotes the essential supremum of L^i under Q^i . Consider the optimizations

$$\hat{Z}_* = \min_{P^i \in \hat{\mathcal{U}}^i, i=1, \dots, m} Z(P^1, \dots, P^m) \quad \text{and} \quad \hat{Z}^* = \max_{P^i \in \hat{\mathcal{U}}^i, i=1, \dots, m} Z(P^1, \dots, P^m) \quad (8)$$

where each $\hat{\mathcal{U}}^i$ contains discrete distributions supported on $\{y_1^i, \dots, y_{n^i}^i\}$, defined in one of the two cases below. For each case, we also make additional assumptions as follows:

1. Set

$$\hat{\mathcal{U}}^i = \{P^i : E_{P^i}[f_l^i(X^i)] \leq \mu_l^i, l = 1, \dots, s^i, \text{ supp } P^i \subset \{y_1^i, \dots, y_{n^i}^i\}\} \quad (9)$$

Moreover, assume that P_0^i satisfies $E_{P_0^i}|f_l^i(X^i)| < \infty$ and $E_{P_0^i}[f_l^i(X^i)] < \mu_l^i$ for all $l = 1, \dots, s^i$.

2. The distribution Q^i is chosen such that P_b^i is absolutely continuous with respect to Q^i , and we denote $L_b^i = dP_b^i/dQ^i$. Set

$$\hat{\mathcal{U}}^i = \{P^i : d_\phi(P^i, \hat{P}_b^i) \leq \eta^i\} \quad (10)$$

where \hat{P}_b^i is defined as

$$\hat{P}_b^i = \sum_{j=1}^{n^i} \frac{L_b^i(y_j^i)}{\sum_{r=1}^{n^i} L_b^i(y_r^i)} \delta(y_j^i)$$

with $\delta(y)$ denoting the delta measure at y . Moreover, assume P_0^i satisfies $E_{P_0^i}|\phi(dP_0^i/dP_b^i)| < \infty$ and $d_\phi(P_0^i, P_b^i) < \eta^i$. Additionally, assume $\phi(\cdot)$ satisfies the continuity condition $|\phi(t(1 + \lambda)) - \phi(t)| \leq |\phi(t)|\kappa_1(\lambda) + \kappa_2(\lambda)$ for any $t \geq 0$ and λ in a fixed neighborhood of 0, where $\kappa_1(\cdot)$ and $\kappa_2(\cdot)$ are two functions such that $\kappa_1(\lambda) = O(\lambda)$ and $\kappa_2(\lambda) = O(\lambda)$ as $\lambda \rightarrow 0$.

Then we have

$$\hat{Z}_* \leq Z_0 + O_p\left(\frac{1}{\sqrt{n}}\right) \leq \hat{Z}^* \quad (11)$$

Here $O_p(1/\sqrt{n})$ is an error term e_n that is of stochastic order $1/\sqrt{n}$, i.e., for any $0 < \epsilon < 1$, there exist $M, N > 0$ such that $P(|\sqrt{n}e_n| < M) > 1 - \epsilon$ for any $n > N$. Theorem 1 is proved in Appendix EC.1. We have a few immediate remarks:

1. Optimizations (8) are the sample counterparts of the original worst-case optimizations (1) with uncertainty sets given by (5) or (7), which optimize discrete distributions over support points that are sampled from generating distributions Q^i 's. Theorem 1 guarantees that, if the original worst-case optimizations give valid covering bounds for the true performance measure (in the spirit of Proposition 1), then so are the sample counterparts, up to an error $O_p(1/\sqrt{n})$ where n denotes the order of the sample size used to construct the sets of support points. The constant implicit in this $O_p(1/\sqrt{n})$ error depends on the sensitivity of Z with respect to the input distributions, as well as the discrepancies between the true input distributions and the support-generating distributions.

2. The condition $\|L^i\|_\infty < \infty$ implies that Q^i has a tail at least as heavy as P_0^i . In practice, the tail of the true distribution P_0^i is not exactly known a priori. This means that it is safer to sample the support points from a heavy-tailed distribution. Additionally, in the case of ϕ -divergence, the generating distribution should also support the baseline. One easy choice is to merely use the baseline as the generating distribution.

3. The conditions $E_{P_0^i}[f_l^i(X^i)] < \mu_l^i$ and $d_\phi(P_0^i, P_b^i) < \eta^i$ state that $E_{P_0^i}[f_l^i(X^i)]$ and $d_\phi(P_0^i, P_b^i)$ are in the interior of $\{(z_1, \dots, z_{s^i}) : z_l \leq \mu_l^i, l = 1, \dots, s^i\}$ and $\{z : z \leq \eta^i\}$ respectively. These conditions guarantee that P_0^i projected on a sample approximation of the support is asymptotically feasible for (8), which helps lead to the guarantee (11). In general, the closer P_0^i is to the boundary of the uncertainty set, i.e., the smaller the values of $\mu_l^i - E_{P_0^i}[f_l^i(X^i)]$ and $\eta^i - d_\phi(P_0^i, P_b^i)$, the larger the sample size is needed for the asymptotic behavior in (11) to kick in, a fact that is not revealed explicitly in Theorem 1. One way to control this required sample size is to expand the uncertainty set by a small margin, say $\epsilon > 0$, i.e., use $E_{P^i}[f_l^i(X^i)] \leq \mu_l^i + \epsilon$ and $d_\phi(P^i, P_b^i) \leq \eta^i + \epsilon$, in (9) and (10). Note that, in the case of moment equality constraint, say $E_{P^i}[f_l^i(X^i)] = \mu_l^i$, one does have to deliberately relax the constraint to $\mu_l^i - \epsilon \leq E_{P^i}[f_l^i(X^i)] \leq \mu_l^i + \epsilon$ for the interior-point conditions to hold.

4. The continuity assumption imposed on $\phi(\cdot)$ in Case 2 is satisfied by many common choices, including KL, (modified) χ^2 -distance, and Burg entropy (see the definitions in Ben-Tal et al. (2013)).

5. As n^i increases, the sampled uncertainty set $\hat{\mathcal{U}}^i$ enlarges as it contains distributions supported on more values. As a result, \hat{Z}_* becomes smaller and \hat{Z}^* larger as n^i increases. Moreover, since $\hat{\mathcal{U}}^i \subset \mathcal{U}^i$, we have $\hat{Z}_* \geq Z_*$ and $\hat{Z}^* \leq Z^*$. This means that as the generated support size increases, the interval $[\hat{Z}_*, \hat{Z}^*]$ progressively widens and is always contained by the interval $[Z_*, Z^*]$.

3.2. Statistical Implications

We further discuss the statistical guarantees implied from Section 3.1. First, a probabilistic analog of Proposition 1 is:

PROPOSITION 2. *Suppose \mathcal{U}^i contains the true distribution P_0^i for all i with confidence $1 - \alpha$, i.e. $\mathbb{P}(\mathcal{U}^i \ni P_0^i \text{ for all } i = 1, \dots, m) \geq 1 - \alpha$, then $\mathbb{P}(Z_* \leq Z_0 \leq Z^*) \geq 1 - \alpha$, where \mathbb{P} denotes the probability generated from a combination of data and prior belief.*

Proposition 2 follows immediately from Proposition 1. In the frequentist framework, \mathbb{P} refers to the probability generated from data. However, Proposition 2 can also be cast in a Bayesian framework, in which \mathbb{P} can represent the prior (e.g., from expert opinion) or the posterior belief.

Proposition 2 reconciles with the established framework in distributionally robust optimization that the uncertainty set \mathcal{U}^i should be chosen as a confidence set for the true distribution, in order to provide a guarantee for the coverage probability on the true objective, in the case that \mathbb{P} represents the generation of data under a true model. Some strategies for constructing confidence sets are:

1. For moment constraint $E_{P^i}[f_l^i(X^i)] \leq \mu_l^i$, one can choose μ_l^i as the upper confidence bound of the moment.
2. For the sup-norm constraint in (6), supposing that P^i is continuous, η^i chosen as the $(1 - \alpha)$ -quantile of $\sup_{x \in [0,1]} B(t)/\sqrt{n^i}$, where $B(t)$ is a standard Brownian bridge, gives an approximate $(1 - \alpha)$ confidence region. This follows from the limiting distribution of the Kolmogorov-Smirnov statistic (see, e.g., Bertsimas et al. (2014)). This calibration becomes conservative (but still correct) when P^i is discrete, and one could use the bootstrap as a remedy. Note that the Kolmogorov-Smirnov-based confidence region is crude for the tail in that it can include a wide range of tail

behaviors, and thus is not recommended if the performance measure of interest is sensitive to the tail.

3. For the ϕ -divergence-based constraint in (7), under the assumption that P^i has finite support of size r^i , Ben-Tal et al. (2013) proposes using $\eta^i = (\phi''(1)/(2n^i))\chi_{r^i-1,1-\alpha}^2$ in the case P_b^i is taken as the empirical distribution, where $\chi_{r^i-1,1-\alpha}^2$ is the $(1-\alpha)$ -quantile of a χ^2 -distribution with degree of freedom $r^i - 1$. This leads to an approximate $(1-\alpha)$ confidence region by using the asymptotics of goodness-of-fit statistics (Pardo (2005)). The resulting region from this approach, however, can be conservative as the involved degree of freedom can be large. Recent works such as Lam and Zhou (2015), Duchi et al. (2016), Lam (2016a) investigate the tightening of divergence-based regions and extend their use to continuous data using the empirical likelihood theory. This theory can also potentially shed insights on the (second-order) accuracies achieved using different divergences (Owen (2001)). Other alternatives include using the Wasserstein distance; see, e.g., Esfahani and Kuhn (2015), Blanchet and Murthy (2016), Gao and Kleywegt (2016) for these developments and the involved ball-size calibration methods.

When discretization is applied, the probabilistic analog of Theorem 1 is:

THEOREM 2. *Suppose all assumptions in Theorem 1 are in place except that $E_{P_0^i}[f_l^i(X^i)] < \mu_l^i$ or $d_\phi(P_0^i, P_b^i) < \eta^i$ now holds true jointly for all i with confidence $1 - \alpha$ under \mathbb{P} . Then $\mathbb{P}(\hat{Z}_* \leq Z_0 + O_p(1/\sqrt{n}) \leq \hat{Z}^*) \geq 1 - \alpha$.*

Theorem 2 follows immediately from Theorem 1. Like before, Theorem 2 translates (1), whose input models can be continuously represented, to (8) that is imposed over discrete distributions, by paying a small price of error. In the next section we discuss our algorithm run over discrete distributions and point out clearly why the discretization is necessary when the input distributions are continuous.

We close this section with two cautionary remarks. First, while our discretization strategy works for problems involving independent low-dimensional input distributions (which occur often in stochastic simulation), high-dimensional joint dependent models may greatly inflate the constant

implicit in the error term, and we do not advise using our strategy in such settings. Second, in general, the finer the discretization scale (i.e., the more generated support points), the higher is the decision space dimension for the resulting optimization problem, and there is a tradeoff on the discretization scale between the approximation error and the optimization effort. Obviously, when the input model is finite discrete, the sampling step depicted in Theorems 1 and 2 is unnecessary, and our subsequent results regarding the algorithm applies readily to this case.

4. Gradient Estimation on Probability Simplices via a Nonparametric Likelihood Ratio Method

Since we work in the discrete space, for simplicity we denote $\mathbf{p}^i = (p_j^i)_{j=1,\dots,n^i} \in \mathbb{R}^{n^i}$ as the vector of probability weights for the discretized input model i . This probability vector is understood to apply on the support points $\{y_1^i, \dots, y_{n^i}^i\}$. Moreover, let $\mathbf{p} = \text{vec}(\mathbf{p}^i : i = 1, \dots, m) \in \mathbb{R}^N$ where vec denotes a concatenation of the vectors \mathbf{p}^i 's as a single vector, and $N = \sum_{i=1}^m n^i$. We denote $\mathcal{P}_i = \{(p_1, \dots, p_l) \in \mathbb{R}^l : \sum_{j=1}^l p_j = 1, p_j \geq 0, j = 1, \dots, l\}$ as the l -dimensional probability simplex. Hence $\mathbf{p}^i \in \mathcal{P}_{n^i}$. For convenience, let $\mathcal{P} = \prod_{i=1}^m \mathcal{P}_{n^i}$, so that $\mathbf{p} \in \mathcal{P}$. The performance measure in (8) can be written as $Z(\mathbf{p})$. Furthermore, denote $T = \max_{i=1,\dots,m} T^i$ as the maximum length of replications among all input models. We also write $\mathbf{X} = (\mathbf{X}^1, \dots, \mathbf{X}^m)$ and $h(\mathbf{X}) = h(\mathbf{X}^1, \dots, \mathbf{X}^m)$ for simplicity. Recall that $I(E)$ denotes the indicator function for the event E . In the rest of this paper, $'$ denotes transpose, and $\|\mathbf{x}\|$ denotes the Euclidean norm of a vector \mathbf{x} . We also write $\text{Var}_{\mathbf{p}}(\cdot)$ as the variance under the input distribution \mathbf{p} . Inequalities for vectors are defined component-wise.

We shall present an iterative simulation-based scheme for optimizing (8). The first step is to design a method to extract the gradient information of $Z(\mathbf{p})$. Note that the standard gradient of $Z(\mathbf{p})$, which we denote as $\nabla Z(\mathbf{p})$, obtained through differentiation of $Z(\mathbf{p})$, may not lead to any simulable object. This is because an arbitrary perturbation of \mathbf{p} may shoot out from the set of probability simplices, and the resulting gradient will be a high-degree polynomial in \mathbf{p} that may have no probabilistic interpretation and thus is not amenable to simulation-based estimation.

We address this issue by considering the set of perturbations within the simplices. Our approach resembles the Gateaux derivative on a functional of probability distribution (Serfling (2009)) as

follows. Given any \mathbf{p}^i , define a mixture distribution $(1 - \epsilon)\mathbf{p}^i + \epsilon\mathbf{1}_j^i$, where $\mathbf{1}_j^i$ represents a point mass on y_j^i , i.e. $\mathbf{1}_j^i = (0, 0, \dots, 1, \dots, 0) \in \mathcal{P}_{n^i}$ and 1 is at the j -th coordinate. The number $0 \leq \epsilon \leq 1$ is the mixture parameter. When $\epsilon = 0$, this reduces to the given distribution \mathbf{p}^i . We treat ϵ as a parameter and differentiate $Z(\mathbf{p}^1, \dots, \mathbf{p}^{i-1}, (1 - \epsilon)\mathbf{p}^i + \epsilon\mathbf{1}_j^i, \mathbf{p}^{i+1}, \dots, \mathbf{p}^m)$ with respect to ϵ for each i, j .

More precisely, let

$$\psi_j^i(\mathbf{p}) = \frac{d}{d\epsilon} Z(\mathbf{p}^1, \dots, \mathbf{p}^{i-1}, (1 - \epsilon)\mathbf{p}^i + \epsilon\mathbf{1}_j^i, \mathbf{p}^{i+1}, \dots, \mathbf{p}^m) \Big|_{\epsilon=0}$$

Denote $\psi^i(\mathbf{p}) = (\psi_j^i(\mathbf{p}))_{j=1, \dots, n^i} \in \mathbb{R}^{n^i}$, and $\psi(\mathbf{p}) = \text{vec}(\psi^i(\mathbf{p}) : i = 1, \dots, m) \in \mathbb{R}^N$. We show that ψ possesses the following two properties:

THEOREM 3. *Given $\mathbf{p} \in \mathcal{P}$ such that $\mathbf{p} > \mathbf{0}$, we have:*

1.

$$\nabla Z(\mathbf{p})'(\mathbf{q} - \mathbf{p}) = \sum_{i=1}^m \nabla^i Z(\mathbf{p})'(\mathbf{q}^i - \mathbf{p}^i) = \sum_{i=1}^m \psi^i(\mathbf{p})'(\mathbf{q}^i - \mathbf{p}^i) = \psi(\mathbf{p})'(\mathbf{q} - \mathbf{p}) \quad (12)$$

for any $\mathbf{q}^i \in \mathcal{P}_{n^i}$ and $\mathbf{q} = \text{vec}(\mathbf{q}^i : i = 1, \dots, m)$, where $\nabla^i Z(\mathbf{p}) \in \mathbb{R}^{n^i}$ is the gradient of Z taken with respect to \mathbf{p}^i .

2.

$$\psi_j^i(\mathbf{p}) = E_{\mathbf{p}}[h(\mathbf{X})s_j^i(\mathbf{X}^i)] \quad (13)$$

where $s_j^i(\cdot)$ is defined as

$$s_j^i(\mathbf{x}^i) = \sum_{t=1}^{T^i} \frac{I(x_t^i = y_j^i)}{p_j^i} - T^i \quad (14)$$

for $\mathbf{x}^i = (x_1^i, \dots, x_{T^i}^i) \in \mathbb{R}^{T^i}$.

The proof of Theorem 3 is in Appendix EC.1. The first property above states that $\psi(\mathbf{p})$ and $\nabla Z(\mathbf{p})$ are identical when viewed as directional derivatives, as long as the direction lies within \mathcal{P} . Since the feasible region of optimizations (8) lies in \mathcal{P} , it suffices to focus on $\psi(\mathbf{p})$. The second property above states that $\psi(\mathbf{p})$ can be estimated unbiasedly in a way similar to the classical likelihood ratio method (Glynn (1990), Reiman and Weiss (1989)), with $s_j^i(\cdot)$ playing the role of

the score function. Since this representation holds without assuming any specific parametric form for \mathbf{p} , we view it as a nonparametric version of the likelihood ratio method.

From (13), an unbiased estimator for $\psi_j^i(\mathbf{p})$ using a single simulation run is $(h(\mathbf{X})s_j^i(\mathbf{X}^i))_{i=1,\dots,m}$, where $\mathbf{X} = (\mathbf{X}^1, \dots, \mathbf{X}^m)$ is the sample path. The following provides a bound on the variance of this estimator (See Appendix EC.1 for proof):

LEMMA 1. *Assume $h(\mathbf{X})$ is bounded a.s., i.e. $|h(\mathbf{X})| \leq M$ for some $M > 0$, and that $\mathbf{p} > \mathbf{0}$. Each sample for estimating $\psi_j^i(\mathbf{p})$, given by $h(\mathbf{X})s_j^i(\mathbf{X}^i)$ using one sample path of \mathbf{X} , possesses a variance bounded from above by $M^2 T^i (1 - p_j^i) / p_j^i$.*

The function $\psi(\mathbf{p})$ derived via the above Gateaux derivative framework can be interpreted as a discrete version of the so-called influence function in robust statistics (Hampel (1974), Hampel et al. (2011)), which is commonly used to approximate the first order effect on a given statistics due to contamination of data. In general, the gradient represented by the influence function is defined as an operator on the domain of the random object distributed under \mathbf{p} . Thus, in the continuous case, this object has an infinite-dimensional domain and can be difficult to compute and encode. This is the main reason why we seek for a discretization in the first place.

5. Frank-Wolfe Stochastic Approximation (FWSA)

With the implementable form of the gradient $\psi(\mathbf{p})$ described in Section 4, we design a stochastic nonlinear programming technique to solve (8). We choose to use the Frank-Wolfe method because, for the types of $\hat{\mathcal{U}}^i$ we consider in Section 3, effective routines exist for solving the induced linearized subproblems.

5.1. Description of the Algorithm

For convenience denote $\hat{\mathcal{U}} = \prod_{i=1}^{n^i} \hat{\mathcal{U}}^i$. We focus on the choices of \mathcal{U}^i depicted in Section 2, which are all convex and consequently $\hat{\mathcal{U}}^i$ and also $\hat{\mathcal{U}}$ are convex.

FWSA works as follows. To avoid repetition we focus only on the minimization formulation in (1). First, pretending that $\nabla Z(\mathbf{p})$ can be computed exactly, it iteratively updates a solution sequence $\mathbf{p}_1, \mathbf{p}_2, \dots$ by, given a current solution \mathbf{p}_k , solving

$$\min_{\mathbf{p} \in \hat{\mathcal{U}}} \nabla Z(\mathbf{p}_k)'(\mathbf{p} - \mathbf{p}_k) \quad (15)$$

Let the optimal solution to (15) be \mathbf{q}_k . The quantity $\mathbf{q}_k - \mathbf{p}_k$ gives a feasible minimization direction starting from \mathbf{p}_k (recall that $\hat{\mathcal{U}}$ is convex). This is then used to update \mathbf{p}_k to \mathbf{p}_{k+1} via $\mathbf{p}_{k+1} = \mathbf{p}_k + \epsilon_k(\mathbf{q}_k - \mathbf{p}_k)$ for some step size ϵ_k . This expression can be rewritten as $\mathbf{p}_{k+1} = (1 - \epsilon_k)\mathbf{p}_k + \epsilon_k\mathbf{q}_k$, which can be interpreted as a mixture between the distributions \mathbf{p}_k and \mathbf{q}_k .

When $\nabla Z(\mathbf{p}_k)$ is not exactly known, one can replace it by an empirical counterpart. Theorem 3 suggests that we can replace $\nabla Z(\mathbf{p}_k)$ by $\hat{\psi}(\mathbf{p}_k)$, and so the empirical counterpart of (15) is

$$\min_{\mathbf{p} \in \hat{\mathcal{U}}} \hat{\psi}(\mathbf{p}_k)'(\mathbf{p} - \mathbf{p}_k) \quad (16)$$

where $\hat{\psi}(\mathbf{p}_k)$ is an estimator of $\psi(\mathbf{p}_k)$ using a sample size R_k . Note that all components of $\hat{\psi}(\mathbf{p}_k)$ can be obtained from these R_k sample paths simultaneously. Letting $\hat{\mathbf{q}}_k$ be the optimal solution to (16), the update rule will be $\mathbf{p}_{k+1} = (1 - \epsilon_k)\mathbf{p}_k + \epsilon_k\hat{\mathbf{q}}_k$ for some step size ϵ_k . The sample size R_k at each step needs to grow suitably to compensate for the bias introduced in solving (16). All these are summarized in Procedure 1.

5.2. Solving the Subproblem

By (12) and the separability of uncertainty set $\hat{\mathcal{U}} = \prod_{i=1}^m \hat{\mathcal{U}}^i$, the subproblem at each iteration can be written as

$$\min_{\mathbf{q} \in \hat{\mathcal{U}}} \sum_{i=1}^m \hat{\psi}^i(\mathbf{p})'(\mathbf{q}^i - \mathbf{p}^i) = \sum_{i=1}^m \min_{\mathbf{q}^i \in \hat{\mathcal{U}}^i} \hat{\psi}^i(\mathbf{p})'(\mathbf{q}^i - \mathbf{p}^i) \quad (17)$$

where $\hat{\psi}^i(\mathbf{p}) = (\hat{\psi}_j^i(\mathbf{p}))_{j=1, \dots, n^i}$ is the empirical counterpart of $\psi^i(\mathbf{p})$ obtained in Algorithm 1. Hence (17) can be solved by m separate convex programs. The update step follows by taking $\mathbf{p}_{k+1} = \text{vec}(\mathbf{p}_{k+1}^i : i = 1, \dots, m)$, where $\mathbf{p}_{k+1}^i = (1 - \epsilon_k)\mathbf{p}_k^i + \epsilon_k\hat{\mathbf{q}}_k^i$ and $\hat{\mathbf{q}}_k^i$ is the solution to the i -th separate program.

Algorithm 1 FWSA for solving (1)

Initialization: $\mathbf{p}_1 \in \mathcal{P}$ where $\mathbf{p}_1 > \mathbf{0}$.

Input: Step size sequence ϵ_k , sample size sequence R_k , $k = 1, 2, \dots$.

Procedure: For each iteration $k = 1, 2, \dots$, given \mathbf{p}_k :

1. Repeat R_k times: Compute

$$h(\mathbf{X})s_j^i(\mathbf{X}^i) \quad \text{for all } i = 1, \dots, m$$

using one sample path $\mathbf{X} = (\mathbf{X}^1, \dots, \mathbf{X}^m)$, where $s_j^i(\mathbf{X}^i) = \sum_{t=1}^{T^i} I(X_t^i = y_j^i)/p_j^i - T^i$ for $j = 1, \dots, n^i$ and $i = 1, \dots, m$. Call these R_k i.i.d. replications $\zeta_j^i(r)$, for $j = 1, \dots, n^i$, $i = 1, \dots, m$, $r = 1, \dots, R_k$.

2. Estimate $\psi(\mathbf{p}_k)$ by

$$\hat{\psi}(\mathbf{p}_k) = (\hat{\psi}_j^i(\mathbf{p}_k))_{i=1, \dots, m, j=1, \dots, n^i} = \left(\frac{1}{R_k} \sum_{r=1}^{R_k} \zeta_j^i(r) \right)_{i=1, \dots, m, j=1, \dots, n^i}.$$

3. Solve $\hat{\mathbf{q}}_k \in \operatorname{argmin}_{\mathbf{p} \in \hat{\mathcal{U}}} \hat{\psi}(\mathbf{p}_k)'(\mathbf{p} - \mathbf{p}_k)$.

4. Update $\mathbf{p}_{k+1} = (1 - \epsilon_k)\mathbf{p}_k + \epsilon_k \hat{\mathbf{q}}_k$.

The separate programs in (17) can be efficiently solved for the uncertainty sets considered in Section 3. To facilitate discussion, we denote a generic form of each separate program in (17) as

$$\min_{\mathbf{p}^i \in \hat{\mathcal{U}}^i} \boldsymbol{\xi}' \mathbf{p}^i \tag{18}$$

for an arbitrary vector $\boldsymbol{\xi} = (\xi_j)_{j=1, \dots, n^i} \in \mathbb{R}^{n^i}$.

Case 1 in Theorem 1: Moment and support constraints. Consider $\hat{\mathcal{U}}^i = \{\mathbf{p}^i \in \mathcal{P}^{n^i} : \mathbf{f}_l^{i\prime} \mathbf{p}^i \leq \mu_l^i, l = 1, \dots, s^i\}$ where $\mathbf{f}_l^i = (f_l(y_j^i))_{j=1, \dots, n^i} \in \mathbb{R}^{n^i}$. Then (18) is a linear program.

Case 2 in Theorem 1: ϕ -divergence neighborhood. Consider

$$\hat{\mathcal{U}}^i = \{\mathbf{p}^i \in \mathcal{P}^{n^i} : d_\phi(\mathbf{p}^i, \mathbf{p}_b^i) \leq \eta^i\} \tag{19}$$

where $\mathbf{p}_b^i = (p_{b,j}^i)_{j=1,\dots,n^i} \in \mathcal{P}^{n^i}$ and $d_\phi(\mathbf{p}^i, \mathbf{p}_b^i) = \sum_{j=1}^{n^i} p_{b,j}^i \phi(p_j^i / p_{b,j}^i)$. We have:

PROPOSITION 3. Consider (18) with $\hat{\mathcal{U}}^i$ presented in (19), where $\mathbf{p}_b^i > \mathbf{0}$. Let $\phi^*(t) = \sup_{x \geq 0} \{tx - \phi(x)\}$ be the conjugate function of ϕ , and define $0\phi^*(s/0) = 0$ if $s \leq 0$ and $0\phi^*(s/0) = +\infty$ if $s > 0$.

Solve the program

$$(\alpha^*, \lambda^*) \in \operatorname{argmax}_{\alpha \geq 0, \lambda \in \mathbb{R}} \left\{ -\alpha \sum_{j=1}^{n^i} p_{b,j}^i \phi^* \left(-\frac{\xi_j + \lambda}{\alpha} \right) - \alpha \eta^i - \lambda \right\} \quad (20)$$

An optimal solution $\mathbf{q}^i = (q_j^i)_{j=1,\dots,n^i}$ for (18) is

1. If $\alpha^* > 0$, then

$$q_j^i = p_{b,j}^i \cdot \operatorname{argmax}_{r \geq 0} \left\{ -\frac{\xi_j + \lambda^*}{\alpha^*} r - \phi(r) \right\} \quad (21)$$

2. If $\alpha^* = 0$, then

$$q_j^i = \begin{cases} \frac{p_{b,j}^i}{\sum_{j \in \mathcal{M}^i} p_{b,j}^i} & \text{for } j \in \mathcal{M}^i \\ 0 & \text{otherwise} \end{cases} \quad (22)$$

where $\mathcal{M}^i = \operatorname{argmin}_j \xi_j$, the set of indices $j \in \{1, \dots, n^i\}$ that have the minimum ξ_j .

Operation (20) involves a two-dimensional convex optimization. Note that both the function ϕ^* and the solution to the n^i one-dimensional maximization (21) have closed-form expressions for all common ϕ -divergence (Pardo (2005)). The proof of Proposition 3 follows closely from Ben-Tal et al. (2013) and is left to Appendix EC.1.

In the special case where $\phi = x \log x - x + 1$, i.e. KL divergence, the solution scheme can be simplified to a one-dimensional root-finding problem. More precisely, we have

PROPOSITION 4. Consider (18) with $\hat{\mathcal{U}}^i$ presented in (19), where $\phi(x) = x \log x - x + 1$ and $\mathbf{p}_b^i > \mathbf{0}$.

Denote $\mathcal{M}^i = \operatorname{argmin}_j \xi_j$ as in Proposition 3. An optimal solution $\mathbf{q}^i = (q_j^i)_{j=1,\dots,n^i}$ for (18) is:

1. If $-\log \sum_{j \in \mathcal{M}^i} p_{b,j}^i \leq \eta^i$, then

$$q_j^i = \begin{cases} \frac{p_{b,j}^i}{\sum_{j \in \mathcal{M}^i} p_{b,j}^i} & \text{for } j \in \mathcal{M}^i \\ 0 & \text{otherwise} \end{cases} \quad (23)$$

2. If $-\log \sum_{j \in \mathcal{M}^i} p_{b,j}^i > \eta^i$, then

$$q_j^i = \frac{p_{b,j}^i e^{\beta \xi_j}}{\sum_{j=1}^{n^i} p_{b,j}^i e^{\beta \xi_j}} \quad (24)$$

for all j , where $\beta < 0$ satisfies

$$\beta \varphi_{\xi}^{i'}(\beta) - \varphi_{\xi}^i(\beta) = \eta^i \quad (25)$$

Here $\varphi_{\xi}^i(\beta) = \log \sum_{j=1}^{n^i} p_{b,j}^i e^{\beta \xi_j}$ is the logarithmic moment generating function of ξ under \mathbf{p}_b^i .

The proof of Proposition 4 follows from techniques in, e.g., Hansen and Sargent (2008), and is left to Appendix EC.1.

6. Theoretical Guarantees of FWSA

This section shows the convergence properties of our proposed FWSA. We first present results on almost sure convergence, followed by a local convergence rate analysis. Throughout our analysis we assume that the subproblem at any iteration can be solved using deterministic optimization routine to a negligible error.

6.1. Almost Sure Convergence

An important object that we will use in our analysis is the so-called Frank-Wolfe (FW) gap (Frank and Wolfe (1956)): For any $\tilde{\mathbf{p}} \in \hat{\mathcal{U}}$, let $g(\tilde{\mathbf{p}}) = -\min_{\mathbf{p} \in \hat{\mathcal{U}}} \psi(\tilde{\mathbf{p}})'(\mathbf{p} - \tilde{\mathbf{p}})$, which is the negation of the optimal value of the next subproblem when the current solution is $\tilde{\mathbf{p}}$. Note that $g(\tilde{\mathbf{p}})$ is non-negative for any $\tilde{\mathbf{p}} \in \hat{\mathcal{U}}$, since one can always take $\mathbf{p} = \tilde{\mathbf{p}}$ in the definition of $g(\tilde{\mathbf{p}})$ to get a lower bound 0. In the case of convex objective function, it is well-known that $g(\tilde{\mathbf{p}})$ provides an upper bound of the actual optimality gap (Frank and Wolfe (1956)). However, we shall make no convexity assumption in our subsequent analysis, and will see that $g(\tilde{\mathbf{p}})$ still plays an important role in bounding the local convergence rate of our procedure under the conditions we impose.

Our choices on the step size ϵ_k and sample size per iteration R_k of the procedure are as follows:

ASSUMPTION 1. We choose $\epsilon_k, k = 1, 2, \dots$ that satisfy

$$\sum_{k=1}^{\infty} \epsilon_k = \infty \quad \text{and} \quad \sum_{k=1}^{\infty} \epsilon_k^2 < \infty$$

ASSUMPTION 2. *The sample sizes $R_k, k = 1, 2, \dots$ are chosen such that*

$$\sum_{k=1}^{\infty} \frac{\epsilon_k}{\sqrt{R_k}} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2} < \infty$$

where for convenience we denote $\prod_{j=1}^0 (1 - \epsilon_j)^{-1/2} = 1$.

Note that among all ϵ_k in the form c/k^α for $c > 0$ and $\alpha > 0$, only $\alpha = 1$ satisfies both Assumptions 1 and 2 and avoids a super-polynomial growth in R_k simultaneously (recall that R_k represents the simulation effort expended in iteration k , which can be expensive). To see this, observe that Assumption 1 asserts $\alpha \in (1/2, 1]$. Now, if $\alpha < 1$, then it is easy to see that $\prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2}$ grows faster than any polynomials, so that R_k cannot be polynomial if Assumption 2 needs to hold. On the other hand, when $\alpha = 1$, then $\prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2}$ grows at rate \sqrt{k} and it is legitimate to choose R_k growing at rate k^β with $\beta > 1$.

Assumption 1 is standard in the SA literature. The growing per-iteration sample size in Assumption 2 is needed to compensate for the bias caused by the subproblem in FWSA. Note that in standard SA, a solution update is obtained by moving in the gradient descent direction, and Assumption 1 suffices if this direction is estimated unbiasedly. In FWSA, the subprogram introduces bias on the feasible direction despite the unbiasedness of the gradient estimate. The increasing simulation effort at each iteration is introduced to shrink this bias as the iteration proceeds. We also note that the expression $\prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2}$ in Assumption 2 is imposed to compensate for a potentially increasing estimation variance, due to the form of the gradient estimator depicted in (13) and (14) that possesses p_j^i in the denominator and thus the possibility of having a larger variance as the iteration progresses.

We state our result on almost sure convergence in two parts. The first part only assumes the continuity of $g(\cdot)$. The second part assumes a stronger uniqueness condition on the optimal solution, stated as:

ASSUMPTION 3. *There exists a unique minimizer \mathbf{p}^* for $\min_{\mathbf{p} \in \hat{\mathcal{U}}} Z(\mathbf{p})$. Moreover, $g(\cdot)$ is continuous over $\hat{\mathcal{U}}$ and \mathbf{p}^* is the only feasible solution such that $g(\mathbf{p}^*) = 0$.*

In light of Assumption 3, g plays a similar role as the gradient in unconstrained problems. The condition $g(\mathbf{p}^*) = 0$ in Assumption 3 is a simple implication of the optimality of \mathbf{p}^* (since $g(\mathbf{p}^*) > 0$ would imply the existence of a better solution).

Our convergence result is:

THEOREM 4. *Suppose that $h(\mathbf{X})$ is bounded a.s. and that Assumptions 1-2 hold. We have the following properties on \mathbf{p}_k generated in Algorithm 1 :*

1. *Assume that $g(\cdot)$ is continuous and an optimal solution exists. Then $D(Z(\mathbf{p}_k), \mathcal{Z}^*) \rightarrow 0$ a.s., where $\mathcal{Z}^* = \{Z(\mathbf{p}) : \mathbf{p} \text{ satisfies } g(\mathbf{p}) = 0\}$ and $D(x, A) = \inf_{y \in A} \|x - y\|$ for any point x and set A in the Euclidean space.*
2. *Under Assumption 3, \mathbf{p}_k converge to \mathbf{p}^* a.s..*

Part 1 of Theorem 4 states that the objective value generated by Algorithm 1 will eventually get close to an objective value evaluated at a point where the FW gap is zero. Part 2 strengthens the convergence to the unique optimal solution \mathbf{p}^* under Assumption 3. In practice, this uniqueness condition may not hold, and we propose combining Algorithm 1 with multi-start of the initial solution \mathbf{p}_1 as a remedy. Section 7.1 and Appendix EC.2.1 show some numerical results on this strategy.

6.2. Local Convergence Rate

We impose several additional assumptions. The first is a Lipchitz continuity condition on an optimal solution for the generic subproblem (18), with respect to the coefficients in the objective in a neighborhood of the gradient evaluated at \mathbf{p}^* . Denote $\mathbf{v}(\boldsymbol{\xi})$ as an optimal solution of (18).

ASSUMPTION 4. *We have*

$$\|\mathbf{v}(\boldsymbol{\xi}_1) - \mathbf{v}(\boldsymbol{\xi}_2)\| \leq L \|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|$$

for some $L > 0$, for any $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2 \in \mathcal{N}_\Delta(\psi(\mathbf{p}^*))$, where $\mathcal{N}_\Delta(\psi(\mathbf{p}^*))$ denotes a Euclidean neighborhood of $\psi(\mathbf{p}^*)$ with radius Δ , and \mathbf{p}^* is assumed to be the unique optimal solution for $\min_{\mathbf{p} \in \hat{\mathcal{U}}} Z(\mathbf{p})$.

Next, we denote $\mathbf{q}(\tilde{\mathbf{p}})$ as an optimizer in the definition of the FW gap at $\tilde{\mathbf{p}}$, i.e. $\mathbf{q}(\tilde{\mathbf{p}}) \in \operatorname{argmin}_{\mathbf{p}} \psi(\tilde{\mathbf{p}})'(\mathbf{p} - \tilde{\mathbf{p}})$.

ASSUMPTION 5.

$$g(\mathbf{p}) \geq c \|\psi(\mathbf{p})\| \|\mathbf{q}(\mathbf{p}) - \mathbf{p}\|$$

for any $\mathbf{p} \in \hat{\mathcal{U}}$, where $c > 0$ is a small constant.

ASSUMPTION 6.

$$\|\psi(\mathbf{p})\| > \tau > 0$$

for any $\mathbf{p} \in \hat{\mathcal{U}}$, for some constant τ .

Assumption 5 guarantees that the angle between the descent direction and the gradient must be bounded away from 90° uniformly at any point \mathbf{p} . This assumption has been used in the design and analysis of gradient descent methods for nonlinear programs that are singular (i.e. without assuming the existence of the Hessian matrix; Bertsekas (1999), Proposition 1.3.3).

The non-zero gradient condition in Assumption 6 effectively suggests that a local optimum must occur at the relative boundary of $\hat{\mathcal{U}}$ (i.e. the boundary with respect to the lower-dimensional subspace induced by the probability simplex constraint), which warrants further explanation. Note that the other alternate scenario for local optimality will be that it occurs in the interior region of the feasible set $\hat{\mathcal{U}}$. In the latter scenario, the gradient at the optimal solution is zero. While the convergence analysis can be simplified (and plausibly give a better rate) under this scenario, the statistical implication brought by this scenario is rather pathological. Note that our optimizations are imposed on decision variables that are input probability distributions. As discussed at the end of Section 4, the gradient vector $\psi(\mathbf{p})$ is the influence function for the performance measure $Z(\cdot)$. If the influence function is zero, it is known that a Gaussian limit does not hold in the central limit theorem as the input sample size gets large (where the central limit theorem is on the difference between a simulation driven by empirical distributions and the truth). Instead, a χ^2 -limit occurs (Serfling (2009), Section 6.4.1, Theorem B). Such type of limit is unusual and has never

been reported in simulation analysis. Indeed, in all our experiments, the obtained local optimal solution is always at the boundary. For this reason we impose Assumption 6 rather than a more straightforward zero-gradient type condition.

The following are our main results on convergence rate, first on the FW gap $g(\mathbf{p}_k)$, and then the optimality gap $Z(\mathbf{p}_k) - Z(\mathbf{p}^*)$, in terms of the number of iterations k . Similar to almost sure convergence, we assume here that the deterministic routine for solving the subproblems can be carried out with high precision.

THEOREM 5. *Suppose $|h(\mathbf{X})| \leq M$ for some $M > 0$ and that Assumptions 1-6 hold. Additionally, set*

$$\epsilon_k = \frac{a}{k} \quad \text{and} \quad R_k = b k^\beta$$

when $k > a$, and arbitrary $\epsilon_k < 1$ when $k \leq a$. Given any $0 < \varepsilon < 1$, it holds that, with probability $1 - \varepsilon$, there exists a large enough positive integer k_0 and small enough positive constants ν, ϑ, ϱ such that $0 < g(\mathbf{p}_{k_0}) \leq \nu$, and for $k \geq k_0$,

$$g(\mathbf{p}_k) \leq \frac{A}{k^C} + B \times \begin{cases} \frac{1}{(C-\gamma)k^\gamma} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^{\gamma-C}k^C} & \text{if } \gamma > C \\ \frac{\log((k-1)/(k_0-1))}{k^C} & \text{if } \gamma = C \end{cases} \quad (26)$$

where

$$A = g(\mathbf{p}_{k_0})k_0^C,$$

$$B = \left(1 + \frac{1}{k_0}\right)^C \left(a\varrho + \frac{2a^2\varrho K}{c\tau k_0} \left(\frac{\nu}{c\tau} + L\vartheta\right)\right)$$

and

$$C = a \left(1 - \frac{2KL\vartheta}{c\tau} - \frac{2K\nu}{c^2\tau^2}\right) \quad (27)$$

Here the constants L, c, τ appear in Assumptions 4, 5 and 6 respectively. The sample size power β needs to be chosen such that $\beta > 2\gamma + a + 1$. More precisely, the constants a, b, β that appear in the specification of the algorithm, the other constants $k_0, \vartheta, \varrho, \gamma, K$, and two new constants $\rho > 1$ and $\delta > 0$ are chosen to satisfy Conditions 1-9 listed in Appendix EC.1.

COROLLARY 1. Suppose that all the assumptions are satisfied and all the constants are chosen as indicated in Theorem 5. Then with probability $1 - \varepsilon$, there exists a large enough positive integer k_0 and small enough positive constants ν, ϑ, ϱ such that $0 \leq g(\mathbf{p}_{k_0}) \leq \nu$, and for $k \geq k_0$,

$$Z(\mathbf{p}_k) - Z(\mathbf{p}^*) \leq \frac{D}{k-1} + \frac{E}{(k-1)^C} + F \times \begin{cases} \frac{1}{(C-\gamma)\gamma(k-1)^\gamma} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^\gamma - C(k-1)^\gamma} & \text{if } \gamma > C \\ \frac{\log((k-1)/(k_0-1))}{C(k-1)^C} & \text{if } \gamma = C \end{cases} \quad (28)$$

where

$$D = a^2 K, \quad E = \frac{aA}{C}, \quad F = aB$$

and a, A, B, C, K are the same constants as in Theorem 5.

A quick summary extracted from Theorem 5 and Corollary 1 is the following: Consider the local convergence rate denominated by workload, i.e. the number of simulation replications. To achieve the most efficient rate, approximately speaking, a should be chosen to be $1 + \omega$ and β chosen to be $5 + \zeta + \omega$ for some small $\omega, \zeta > 0$. The local convergence rate is then $O(W^{-1/(6+\zeta+\omega)})$ where W is the total number of simulation replications.

Note that the bounds in Theorem 5 and Corollary 1 are local asymptotic statements since they only hold starting from $k \geq k_0$ and $g(\mathbf{p}_k) \leq \nu$ for some large k_0 and small ν . It should be cautioned that they do not say anything about the behavior of the algorithm before reaching the small neighborhood of \mathbf{p}^* as characterized by $0 \leq g(\mathbf{p}_{k_0}) \leq \nu$. The above summary therefore should be interpreted in the way that, given the algorithm has already run k_0 number of replications and $g(\mathbf{p}_k) \leq \nu$ for a suitably small ν (which occurs with probability 1 by Theorem 4), the convergence rate of $O(W^{-1/(6+\zeta+\omega)})$ for the optimality gap is guaranteed with probability $1 - \varepsilon$ starting from that point.

The summary above is derived based on the following observations:

1. The local convergence rate of the optimality gap, in terms of the number of iterations k , is at best $O(1/k^{C \wedge \gamma \wedge 1})$. This is seen by (28).

2. We now consider the convergence rate in terms of simulation replications. Note that at iteration k , the cumulative number of replications is of order $\sum_{j=1}^k j^\beta \approx k^{\beta+1}$. Thus from Point 1 above, the convergence rate of the optimality gap in terms of replications is of order $1/W^{(C \wedge \gamma \wedge 1)/(\beta+1)}$.

3. The constants C and γ respectively depend on a , the constant factor in the step size, and β , the geometric growth rate of the sample size, as follows:

(a) (27) defines $C = a(1 - 2KL\vartheta/(c\tau) - 2K\nu/(c^2\tau^2))$. For convenience, we let $\omega = 2KL\vartheta/(c\tau) + 2K\nu/(c^2\tau^2)$, and so $C = a(1 - \omega)$.

(b) From Condition 6 in Theorem 5 (shown in Appendix EC.1), we have $\beta = 2\gamma + \rho a + 2 + \zeta$ for some $\zeta > 0$. In other words $\gamma = (\beta - \rho a - \zeta - 2)/2$.

4. Therefore, the convergence rate in terms of replications is $1/W^{((a(1-\omega)) \wedge ((\beta - \rho a - \zeta - 2)/2) \wedge 1)/(\beta+1)}$.

Let us focus on maximizing

$$\frac{(a(1 - \omega)) \wedge ((\beta - \rho a - \zeta - 2)/2) \wedge 1}{\beta + 1} \quad (29)$$

over a and β , whose solution is given by the following lemma:

LEMMA 2. *The maximizer of (29) is given by*

$$a = \frac{1}{1 - \omega}, \quad \beta = \frac{\rho}{1 - \omega} + \zeta + 4$$

and the optimal value is

$$\frac{1}{\rho/(1 - \omega) + \zeta + 5}$$

The proof is in Appendix EC.1. With Lemma 2, let us choose ϑ and ν , and hence ω , to be small. We also choose ρ to be close to 1. (Unfortunately, these choices can lead to a small size of neighborhood around \mathbf{p}^* in which the convergence rate holds.) This gives rise to the approximate choice that $a \approx 1 + \omega$ and $\beta \approx 5 + \zeta + \omega$. The convergence rate is then $O(W^{-1/(6+\zeta+\omega)})$.

We compare our results to some recent work in stochastic FW. Hazan and Luo (2016) showed that to achieve ϵ error in terms of the optimality gap one needs $O(1/\epsilon^{1.5})$ number of calls to the gradient

estimation oracle, when the objective function is strongly convex. Reddi et al. (2016) showed that the number needed increases to $O(1/\epsilon^4)$ for non-convex objectives, and suggested several more sophisticated algorithms to improve the rate. Corollary 1 and our discussion above suggests that we need $O(1/\epsilon^{6+\zeta+\omega})$ sample size, for some small $\zeta, \omega > 0$, a rate that is inferior to the one achieved in Reddi et al. (2016). However, Reddi et al. (2016) has assumed that the gradient estimator is uniformly bounded over the feasible space, a condition known as G -Lipschitz (Theorem 2 in Reddi et al. (2016)), which does not hold in our case due to the presence of p_j^i in the denominator in (14) that gives a potentially increasing estimation variance as the iteration progresses. This complication motivates our sample size and step size sequences depicted in Assumption 2 and the subsequent analysis. On the other hand, if Assumption 4 is relaxed to hold for any $\xi_1, \xi_2 \in \mathbb{R}^N$, it can be seen that by choosing $\beta \approx 3 + \zeta + \omega$ our complexity improves to $O(1/\epsilon^{4+\zeta+\omega})$, which almost matches the one in Reddi et al. (2016) (see Remark EC.1 in Appendix EC.1). However, such a relaxed condition would not hold if the constraints are linear, because the optimal solutions of the subproblems are located at the corner points and will jump from one to the other under perturbation of the objective function.

7. Numerical Experiments

This section describes two sets of numerical experiments. The first set (Section 7.1) studies the performance guarantees from Section 3 regarding our randomized discretization strategy and the tightness of the bounds coming from moment constraints. The second set of experiments (Section 7.2) studies the numerical convergence of FWSA. The appendix provides additional details and results. Unless specified, in all experiments we terminate the FWSA algorithm at iteration k if at least one of the following criteria is met (as an indication that the convergence studied in Section 6 is attained):

- The cumulative simulation replications W_k reaches 5×10^8 , or
- The relative difference between objective value $Z(\mathbf{p}_k)$ and the average of the observed values in 30 previous iterations, $(\sum_{v=1}^{30} Z(\mathbf{p}_{k-v}))/30$, is below 5×10^{-5} , or
- The gradient estimate $\hat{\psi}(\mathbf{p}_k)$ has an l_2 -norm smaller than 1×10^{-3} .

7.1. Performance Bounds for Multiple Continuous and Unbounded Input Models

We use the example of a generic multi-class $M/G/1$ queue where jobs from three distinct classes arrive and are attended to by one server. Such structures are common in service systems such as call-centers. Let $\mathbf{P} = \{P^1, P^2, P^3\}$ represent all the constituent probability measures, where each $P^i = \{P^{i,j}\}$, $i = 1, 2, 3$ with $j = 1$ for interarrival and $j = 2$ for service, denotes the joint measure of the interarrival and service distributions of jobs of class i . The performance measure of interest is the weighed average waiting time:

$$Z(\mathbf{P}) = E_{\mathbf{P}} \left[\sum_{i=1}^3 \left(c_i \frac{1}{T^i} \sum_{t=1}^{T^i} W_t^i \right) \right], \quad (30)$$

where the average is observed up to a (fixed) $T^i = 500$ customers of class i and c_i is the cost assigned to its waiting times. Jobs within each class are served on a first-come-first-served basis. The server uses a fixed priority ordering based on the popular $c\mu$ rule (Kleinrock (1976)), which prioritizes the class on the next serving in decreasing order of the product of c_i and the mean service rate μ^i of class i (as discussed momentarily, the μ^i 's are unknown, so we fix a specific guess throughout this example).

To handle the uncertainty in specifying the interarrival and service time distributions of each class (due to, e.g., the novelty of the service operation with little pre-existing data), we use the uncertainty set based on moment constraints on the P^i as:

$$\mathcal{U} = \prod_i \mathcal{U}^i, \text{ where } \mathcal{U}^i = \{P^i : \underline{\mu}_l^{i,j} \leq E_{P^i}[(X^{i,j})^l] \leq \bar{\mu}_l^{i,j}, l = 1, 2, j = 1, 2\} \quad (31)$$

where the index $l = 1, 2$ represents the first two moments of marginals $P^{i,j}$. This set is motivated from queueing theory that mean system responses could depend on the mean and variance of the input distributions. The moment bounds $\underline{\mu}_l^{i,j}$ and $\bar{\mu}_l^{i,j}$ can be specified from prior or expert opinion. Here, to test the information value with respect to the accuracy of the moments, we specify the bounds from a confidence interval on the corresponding moments calculated from N_s synthetically generated observations for each i, j . For example,

$$\bar{\mu}_l^{i,j} = \hat{\mu}_l^{i,j} + t_{\alpha/2, N_s - 1} \hat{\sigma}_l^{i,j} / \sqrt{N_s},$$

where $t_{\alpha/2, N_s - 1}$ is the $(1 - \alpha/2)$ -quantile of the Student-t distribution with degree of freedom $N_s - 1$, $\hat{\mu}_l^{i,j}$ is the empirical l -th moment and $\hat{\sigma}_l^{i,j}$ is the associated sample standard deviation as observed from the N_s data points. Suppose that the true marginal distribution of interarrival times for each class is exponential with rate 0.5 and the true service distribution of the three classes are exponentials with rates 2.25, 2.0 and 1.75 respectively, to yield an overall traffic intensity of 0.75.

The FWSA algorithm is run by first sampling a discrete approximate support from bivariate independent-marginal lognormal distributions as representative of each P^i with support size $n = 50, 100, 250$ (we assume the support size corresponding to each distribution P^i is all equal to n). Theorem 1 suggests that selecting lognormal distributions is reasonable if the modeler conjectures that the true distributions are light-tailed. Here we set the means and standard deviations of the lognormals to 1. The parameter n should ideally be large to minimize discretization error, but this pays a penalty in the slowness of the FWSA algorithm.

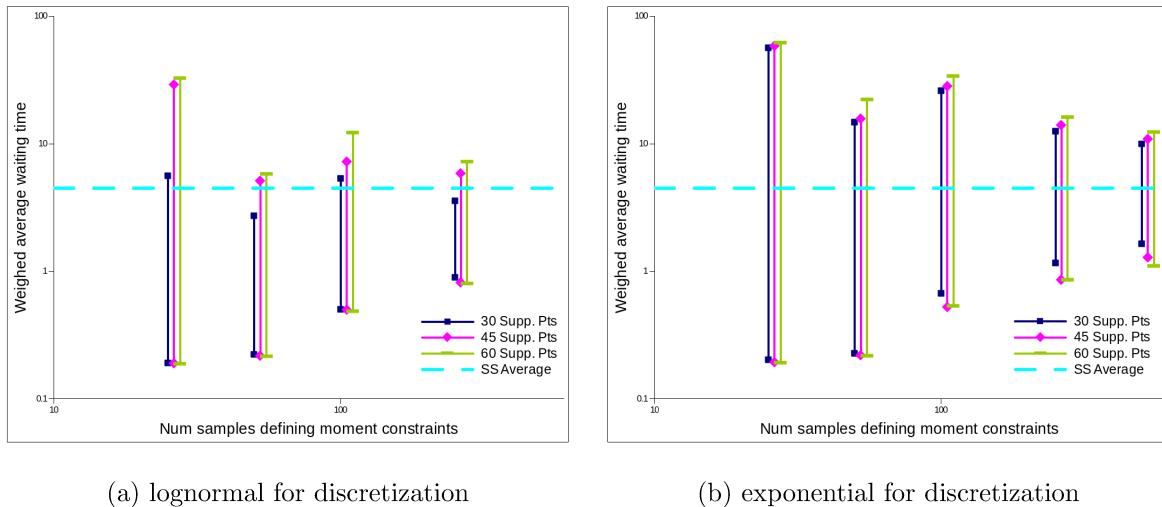


Figure 1 The range from max to min worst-case objectives when N_s and n vary as indicated. The dotted-line indicates the expected steady-state performance under the true distribution.

Figure 1a shows the output of our approach over various n and N_s to illustrate the effect of discretization and the accuracy of moment information on the tightness of our bounds. The true steady-state performance measure of the multiclass $M/M/1$ system, available in explicit form

(Kleinrock (1976)), is indicated as the dotted-line in each plot. The bounds provided by our method are all seen to cover the true performance value when $n \geq 45$. This is predicted by Theorem 1 as the moment constraints are all correctly calibrated (i.e. contain the true moments) in this example. Moreover, as predicted by discussion point 5 in Section 3.1, the obtained intervals widen as n increases, since the expansion of support size enlarges the feasible region. On the other hand, the intervals shrink as N_s increases, since this tightens the moment constraints and consequently reduces the feasible region. The effect of the support size does not appear too sensitive in this example. Thus, taking into account the optimization efficiency, a use of support size of about 45 points appears sufficient.

Figure 1b plots the performance when the supports of the distributions are sampled from the true distributions. The performance trends are similar to Figure 1a. However, the obtained bounds are slightly looser. Note that Theorem 1 guarantees that the obtained bounds under the generated support points cover the truth with high confidence, when the generating distributions satisfy the heavier-tail condition. In this example, both lognormal and exponential distributions (the latter being the truth) satisfy these conditions and lead to correct bounds. On the other hand, the tightness of the bounds, which is not captured in Theorem 1, depends on the size and geometry of the feasible region that is determined by a complex interplay between the choice of the uncertainty set and the support-generating distributions. The feasible region using the true exponential distributions include probability weights that are close to uniform weights (since the moment constraints are calibrated using the same distribution). The region using the lognormal, however, does not contain such weights; in fact, when $N_s = 500$, the resulting optimizations can be infeasible for $n \leq 60$, signaling the need to use more support-generating samples, whereas they are always feasible using the exponential, whose values are shown in the rightmost set of intervals in Figure 1b.

The results above are implemented with an initialization that assigns equal probabilities to the support points. Appendix EC.2.1 shows the results applied on different initializations to provide evidence that the formulation in this example has a unique global optimal solution or similar local optimal solutions.

7.2. Convergence of FWSA and Worst-case Input Distributions

We test the numerical convergence of FWSA. The key parameters in the algorithm are the sample-size growth rate β and the step-size constant a . Varying these two parameters, we empirically test the rate of convergence of the FW gap to zero analyzed in Theorem 5, and the objective function $Z(\mathbf{p}_k)$ to the true optimal value $Z(\mathbf{p}^*)$ analyzed in Corollary 1. We also investigate the magnitude of the optimal objective value and the form of the identified optimal solution.

Here we consider an $M/G/1$ queue where the arrival process is Poisson known with high accuracy to have rate 1. On the other hand, the service time X_t for the t -th customer is uncertain but assumed i.i.d.. A simulation model is being used to estimate the expected long-run average of the waiting times $Z(\mathbf{p}) = E_{\mathbf{p}}[h(\mathbf{X})]$, where

$$h(\mathbf{X}) = \frac{1}{T} \sum_1^T W_t$$

and W_t is the waiting time obtained from Lindley's recursion.

We test our FWSA with a KL-divergence-based uncertainty set for X_t as

$$\hat{\mathcal{U}} = \left\{ \mathbf{p} : \sum_{j=1}^n p_j \log \left(\frac{p_j}{p_{b,j}} \right) \leq \eta \right\} \quad (32)$$

where $\mathbf{p}_b = (p_{b,j})_{j=1,\dots,n}$ is a baseline model chosen to be a discretized mixture of beta distribution given by $0.3 \times \text{Beta}(2, 6) + 0.7 \times \text{Beta}(6, 2)$. The discrete supports are obtained by uniformly discretizing the interval $[0, 1]$ into n points, i.e. $y_j = (j + 1)/n$. The set (32) provides a good testing ground because steady-state analysis allows obtaining an approximate optimal solution directly which serves as a benchmark for verifying the convergence of our FWSA algorithm (see Appendix EC.2.2 for further details of this approximate optimal solution).

Figure 2 captures the performance of our FWSA algorithm as a function of the a and β parameters. Figures 2a–2c plot the (approximate) optimality gap as a function of the cumulative simulation replications W_k for the maximization problem under (32). We set the parameters $\eta = 0.025$, $n = 100$ and $T = 500$. Figures 2a, 2b and 2c provide further insights into the actual observed finite-sample

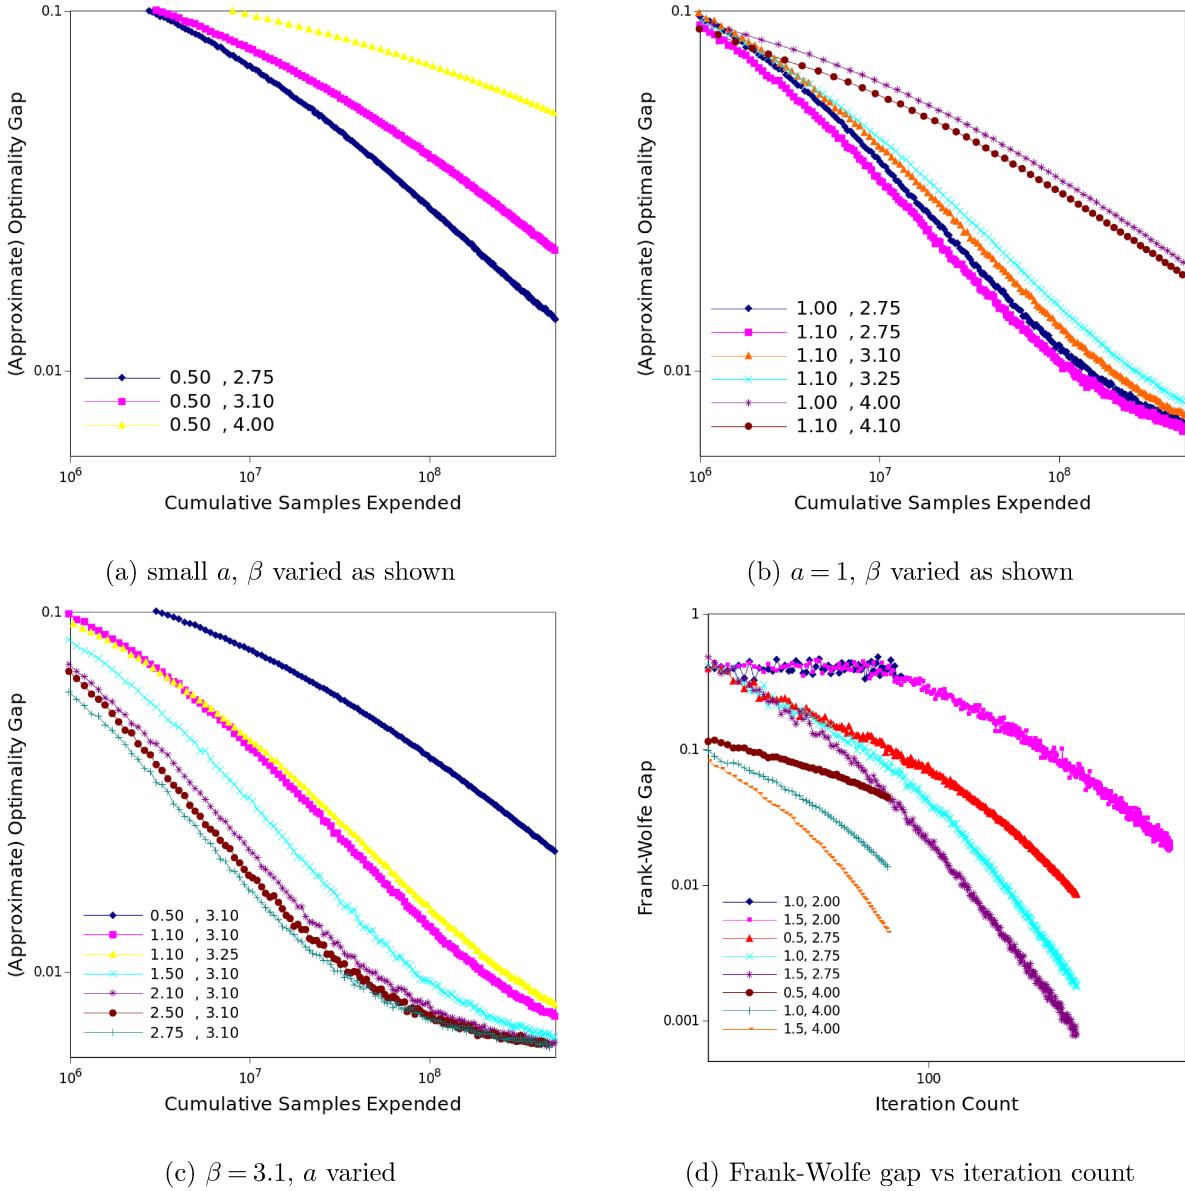


Figure 2 Figs 2a, 2b and 2c plot the optimality gap of the FWSA algorithm for the $M/G/1$ example as function of cumulative simulation samples, under various combinations of step-size parameter a and sample-size growth parameter β . The three figures have the same range of values in both axes (note the log scale).

Fig 2d shows the FW gap as a function of iteration count. All figures provide the legend as a, β .

performance (When interpreting these graphs, note that they are plotted in log-log scale and thus, roughly speaking, the slope of the curve represents the power of the cumulative samples whereas the intercept represents the multiplicative constant in the rate):

- *Fig. 2a v.s. 2b–2c:* Convergence is much slower when $a < 1$ no matter the value of β .

- *Fig. 2b:* For $a > 1$, convergence is again slow if $\beta > 4$.
- *Fig. 2b:* For a slightly greater than 1, the convergence rates are similar for $\beta \in [2.75, 3.25]$ with better performance for the lower end.
- *Fig. 2c:* For $\beta = 3.1$, the rate of convergence generally improves as a increases in the range $[1.10, 2.75]$.
- *Figs. 2a, 2b and 2c:* The approximation Z_∞^* of the true $Z(\mathbf{p}^*)$ (from (SS) in Appendix EC.2.2) has an error of about 0.006 for the chosen T , as observed by the leveling off of all plots around this value as the sampling effort grows.

Figure 2d shows the FW gap as a function of the iteration count. In general, the sample paths with similar β are clustered together, indicating that more effort expended in estimating the gradient at each iterate leads to a faster drop in the FW gap per iteration. Within each cluster, performance is inferior when $a < 1$, consistent with Theorem 5. Since most runs terminate when the criterion on the maximum allowed budget of simulation replications is expended, the end points of the curves indicate that a combination of $a \geq 1$ and a β of around 3 gains the best finite-sample performance in terms of the FW gap. These choices seem to reconcile with the discussion at the end of Section 6.2 when Assumption 4 is relaxed to hold for any $\xi_1, \xi_2 \in \mathbb{R}^N$.

We provide further discussion on the shape of the obtained optimal distributions in Appendix EC.2.3.

8. Conclusion

In this paper we investigated a methodology based on worst-case analysis to quantify input errors in stochastic simulation, by using optimization constraints to represent the partial nonparametric information on the model. The procedure involved a randomized discretization of the support and running FWSA using a gradient estimation technique akin to a nonparametric version of the likelihood ratio or the score function method. We studied the statistical guarantees of the discretization and convergence properties of the proposed FWSA. We also tested our method and verified the theoretical implications on queueing examples.

We suggest several lines of future research. First is the extension of the methodology to dependent models, such as Markovian inputs or more general time series inputs, which would involve new sets of constraints in the optimizations. Second is the design and analysis of other potential alternate numerical procedures and comparisons with the proposed method. Third is the utilization of the proposed worst-case optimizations in various classes of decision-making problems.

Acknowledgments

We thank the Area Editor, the Associate Editor and the three referees for many helpful suggestions that have greatly improved the paper. We gratefully acknowledge support from the National Science Foundation under grants CMMI-1542020, CMMI-1523453 and CAREER CMMI-1653339.

References

Bandi, Chaithanya, Dimitris Bertsimas. 2012. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical programming* **134**(1) 23–70.

Banks, J, JS Carson, BL Nelson, DM Nicol. 2009. *Discrete-Event System Simulation*. 5th ed. Prentice Hall Englewood Cliffs, NJ, USA.

Barton, Russell R, Barry L Nelson, Wei Xie. 2013. Quantifying input uncertainty via simulation confidence intervals. *INFORMS Journal on Computing* **26**(1) 74–87.

Barton, Russell R, Lee W Schruben. 1993. Uniform and bootstrap resampling of empirical distributions. *Proceedings of the 1993 Winter Simulation Conference*. ACM, 503–508.

Barton, Russell R, Lee W Schruben. 2001. Resampling methods for input modeling. *Proceedings of the 2001 Winter Simulation Conference*, vol. 1. IEEE, 372–378.

Bayraksan, Güzin, David K Love. 2015. Data-driven stochastic programming using phi-divergences. *The Operations Research Revolution*. INFORMS, 1–19.

Ben-Tal, Aharon, Dick Den Hertog, Anja De Waegenaere, Bertrand Melenberg, Gijs Rennen. 2013. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science* **59**(2) 341–357.

Ben-Tal, Aharon, Laurent El Ghaoui, Arkadi Nemirovski. 2009. *Robust Optimization*. Princeton University Press.

Bertsekas, Dimitri P. 1999. *Nonlinear Programming*. Athena Scientific.

Bertsimas, Dimitris, David B Brown, Constantine Caramanis. 2011. Theory and applications of robust optimization. *SIAM Review* **53**(3) 464–501.

Bertsimas, Dimitris, Vishal Gupta, Nathan Kallus. 2014. Robust sample average approximation. *Mathematical Programming* 1–66.

Bertsimas, Dimitris, Karthik Natarajan. 2007. A semidefinite optimization approach to the steady-state analysis of queueing systems. *Queueing Systems* **56**(1) 27–39.

Bertsimas, Dimitris, Ioana Popescu. 2005. Optimal inequalities in probability theory: A convex optimization approach. *SIAM Journal on Optimization* **15**(3) 780–804.

Birge, John R, Roger J-B Wets. 1987. Computing bounds for stochastic programming problems by means of a generalized moment problem. *Mathematics of Operations Research* **12**(1) 149–162.

Blanchet, Jose, Yang Kang, Karthyek Murthy. 2016. Robust Wasserstein profile inference and applications to machine learning. *arXiv preprint arXiv:1610.05627* .

Blanchet, Jose, Karthyek RA Murthy. 2016. Quantifying distributional model risk via optimal transport. *arXiv preprint arXiv:1604.01446* .

Blum, Julius R. 1954. Multidimensional stochastic approximation methods. *The Annals of Mathematical Statistics* 737–744.

Boyd, Stephen, Lieven Vandenberghe. 2009. *Convex Optimization*. Cambridge university press.

Buche, Robert, Harold J Kushner. 2002. Rate of convergence for constrained stochastic approximation algorithms. *SIAM Journal on Control and Optimization* **40**(4) 1011–1041.

Canon, MD, CD Cullum. 1968. A tight upper bound on the rate of convergence of Frank-Wolfe algorithm. *SIAM Journal on Control* **6**(4) 509–516.

Cario, Marne C, Barry L Nelson. 1997. Modeling and generating random vectors with arbitrary marginal distributions and correlation matrix. Tech. rep., Citeseer.

Channouf, Nabil, Pierre L'Ecuyer. 2009. Fitting a normal copula for a multivariate distribution with both discrete and continuous marginals. *Proceedings of the Winter Simulation Conference*. IEEE, 352–358.

Cheng, Russell CH, Wayne Holland. 1997. Sensitivity of computer simulation experiments to errors in input data. *Journal of Statistical Computation and Simulation* **57**(1-4) 219–241.

Cheng, Russell CH, Wayne Holland. 1998. Two-point methods for assessing variability in simulation output. *Journal of Statistical Computation Simulation* **60**(3) 183–205.

Cheng, Russell CH, Wayne Holland. 2004. Calculation of confidence intervals for simulation output. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* **14**(4) 344–362.

Chick, Stephen E. 2001. Input distribution selection for simulation experiments: accounting for input uncertainty. *Operations Research* **49**(5) 744–758.

Delage, Erick, Yinyu Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research* **58**(3) 595–612.

Duchi, John C, Peter W Glynn, Hongseok Namkoong. 2016. Statistics of robust optimization: A generalized empirical likelihood approach. *arXiv preprint arXiv:1610.03425* .

Dunn, Joseph C. 1979. Rates of convergence for conditional gradient algorithms near singular and nonsingular extremals. *SIAM Journal on Control and Optimization* **17**(2) 187–211.

Dunn, Joseph C. 1980. Convergence rates for conditional gradient sequences generated by implicit step length rules. *SIAM Journal on Control and Optimization* **18**(5) 473–487.

Dupuis, Paul, Markos A Katsoulakis, Yannis Pantazis, Petr Plecháć. 2016. Path-space information bounds for uncertainty quantification and sensitivity analysis of stochastic dynamics. *SIAM/ASA Journal on Uncertainty Quantification* **4**(1) 80–111.

Esfahani, Peyman Mohajerin, Daniel Kuhn. 2015. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming* 1–52.

Fan, Weiwei, L Jeff Hong, Xiaowei Zhang. 2013. Robust selection of the best. *Proceedings of the 2013 Winter Simulation Conference*. IEEE Press, 868–876.

Frank, Marguerite, Philip Wolfe. 1956. An algorithm for quadratic programming. *Naval Research Logistics Quarterly* **3**(1-2) 95–110.

Freund, Robert M, Paul Grigas. 2014. New analysis and results for the Frank-Wolfe method. *arXiv preprint arXiv:1307.0873v2* .

Fu, Michael C. 1994. Optimization via simulation: A review. *Annals of Operations Research* **53**(1) 199–247.

Gao, Rui, Anton J Kleywegt. 2016. Distributionally robust stochastic optimization with Wasserstein distance. *arXiv preprint arXiv:1604.02199* .

Ghosh, Soumyadip, Shane G Henderson. 2002. Chessboard distributions and random vectors with specified marginals and covariance matrix. *Operations Research* **50**(5) 820–834.

Glasserman, Paul, Xingbo Xu. 2013. Robust portfolio control with stochastic factor dynamics. *Operations Research* **61**(4) 874–893.

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

Glasserman, Paul, Linan Yang. 2016. Bounding wrong-way risk in CVA calculation. *Mathematical Finance* .

Glynn, Peter W. 1990. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM* **33**(10) 75–84.

Goh, Joel, Melvyn Sim. 2010. Distributionally robust optimization and its tractable approximations. *Operations Research* **58**(4-Part-1) 902–917.

Hampel, Frank R. 1974. The influence curve and its role in robust estimation. *Journal of the American Statistical Association* **69**(346) 383–393.

Hampel, Frank R, Elvezio M Ronchetti, Peter J Rousseeuw, Werner A Stahel. 2011. *Robust Statistics: The Approach Based on Influence Functions*, vol. 114. John Wiley & Sons.

Hansen, Lars Peter, Thomas J. Sargent. 2001. Robust control and model uncertainty. *The American Economic Review* **91**(2) pp. 60–66.

Hansen, Lars Peter, Thomas J Sargent. 2008. *Robustness*. Princeton University Press.

Hazan, Elad, Haipeng Luo. 2016. Variance-reduced and projection-free stochastic optimization. *arXiv preprint arXiv:1602.02101* .

Hu, Zhaolin, Jing Cao, L Jeff Hong. 2012. Robust simulation of global warming policies using the DICE model. *Management Science* **58**(12) 2190–2206.

Iyengar, Garud N. 2005. Robust dynamic programming. *Mathematics of Operations Research* **30**(2) 257–280.

Jaggi, Martin. 2013. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*. 427–435.

Jain, A., A. Lim, J. Shanthikumar. 2010. On the optimality of threshold control in queues with model uncertainty. *Queueing Systems* **65** 157–174.

Khintchine, A. Y. 1932. Mathematical theory of a stationary queue. *Matematicheskii Sbornik* **39** 7384.

Kleinrock, Leonard. 1976. *Computer Applications, Volume 2, Queueing Systems*. Wiley-Interscience.

Kushner, H., G. Yin. 2003. *Stochastic Approximation and Recursive Algorithms and Applications*. 2nd ed. Springer-Verlag, New York.

Kushner, Harold J. 1974. Stochastic approximation algorithms for constrained optimization problems. *The Annals of Statistics* 713–723.

Lafond, Jean, Hoi-To Wai, Eric Moulines. 2016. On the online Frank-Wolfe algorithms for convex and non-convex optimizations. *arXiv preprint arXiv:1510.01171* .

Lam, Henry. 2016a. Recovering best statistical guarantees via the empirical divergence-based distributionally robust optimization. *arXiv preprint arXiv:1605.09349* .

Lam, Henry. 2016b. Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research* **41**(4) 1248–1275.

Lam, Henry. 2017. Sensitivity to serial dependency of input processes: A robust approach. *Management Science* .

Lam, Henry, Enlu Zhou. 2015. Quantifying uncertainty in sample average approximation. *Proceedings of the 2015 Winter Simulation Conference*. IEEE Press, 3846–3857.

Lim, Andrew E. B., J. George Shanthikumar. 2007. Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research* **55**(2) 198–214.

Lim, Andrew E. B., J. George Shanthikumar, Thaisiri Watewai. 2011. Robust asset allocation with benchmarked objectives. *Mathematical Finance* **21**(4) 643–679.

Lim, Andrew EB, J George Shanthikumar, ZJ Max Shen. 2006. Model uncertainty, robust optimization and learning. *Tutorials in Operations Research* 66–94.

Luenberger, David G. 1969. *Optimization by Vector Space Methods*. John Wiley & Sons.

Lurie, Philip M, Matthew S Goldberg. 1998. An approximate method for sampling correlated random variables from partially-specified distributions. *Management Science* **44**(2) 203–218.

Nemirovski, Arkadi, Anatoli Juditsky, Guanghui Lan, Alexander Shapiro. 2009. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization* **19**(4) 1574–1609.

Nilim, Arnab, Laurent El Ghaoui. 2005. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research* **53**(5) 780–798.

Owen, Art B. 2001. *Empirical Likelihood*. CRC Press.

Pardo, Leandro. 2005. *Statistical Inference Based on Divergence Measures*. CRC Press.

Pasupathy, Raghu, Sujin Kim. 2011. The stochastic root-finding problem: Overview, solutions, and open questions. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* **21**(3) 19:1–23.

Petersen, Ian R., Matthew R. James, Paul Dupuis. 2000. Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Transactions on Automatic Control* **45**(3) 398–412.

Qian, Pengyu, Zizhuo Wang, Zaiwen Wen. 2015. A composite risk measure framework for decision making under uncertainty. *arXiv preprint arXiv:1501.01126* .

Reddi, Sashank J, Suvrit Sra, Barnabas Poczos, Alex Smola. 2016. Stochastic Frank-Wolfe methods for nonconvex optimization. *arXiv preprint arXiv:1607.08254* .

Reiman, Martin I, Alan Weiss. 1989. Sensitivity analysis for simulations via likelihood ratios. *Operations Research* **37**(5) 830–844.

Rubinstein, Reuven Y. 1986. The score function approach for sensitivity analysis of computer simulation models. *Mathematics and Computers in Simulation* **28**(5) 351–379.

Ryzhov, Ilya O, Boris Defourny, Warren B Powell. 2012. Ranking and selection meets robust optimization. *Proceedings of the 2012 Winter Simulation Conference*. IEEE, 48.

Saltelli, Andrea, Paola Annoni, Ivano Azzini, Francesca Campolongo, Marco Ratto, Stefano Tarantola. 2010. Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Computer Physics Communications* **181**(2) 259–270.

Saltelli, Andrea, Marco Ratto, Terry Andres, Francesca Campolongo, Jessica Cariboni, Debora Gatelli, Michaela Saisana, Stefano Tarantola. 2008. *Global Sensitivity Analysis: The Primer*. Wiley.

Serfling, Robert J. 2009. *Approximation Theorems of Mathematical Statistics*, vol. 162. John Wiley & Sons.

Shafieezadeh-Abadeh, Soroosh, Peyman Mohajerin Esfahani, Daniel Kuhn. 2015. Distributionally robust logistic regression. *Advances in Neural Information Processing Systems*. 1576–1584.

Smith, J. E. 1995. Generalized Chebyshev inequalities: Theory and applications in decision analysis. *Operations Research* **43**(5) 807–825.

Smith, James E. 1993. Moment methods for decision analysis. *Management Science* **39**(3) 340–358.

Song, Eunhye, Barry L Nelson. 2015. Quickly assessing contributions to input uncertainty. *IIE Transactions* **47**(9) 893–909.

Wiesemann, Wolfram, Daniel Kuhn, Melvyn Sim. 2014. Distributionally robust convex optimization. *Operations Research* **62**(6) 1358–1376.

Xie, Wei, Barry L Nelson, Russell R Barton. 2014. A Bayesian framework for quantifying uncertainty in stochastic simulation. *Operations Research* **62**(6) 1439–1452.

Xie, Wei, Barry L Nelson, Russell R Barton. 2015. Statistical uncertainty analysis for stochastic simulation. *Under review in Operations Research*.

Xin, Linwei, David A Goldberg. 2015. Distributionally robust inventory control when demand is a martingale. *arXiv preprint arXiv:1511.09437*.

Xu, Huan, Shie Mannor. 2012. Distributionally robust Markov decision processes. *Mathematics of Operations Research* **37**(2) 288–300.

Zhou, Enlu, Wei Xie. 2015. Simulation optimization when facing input uncertainty. *Proceedings of the 2015 Winter Simulation Conference*. IEEE Press, 3714–3724.

Zouaoui, Faker, James R Wilson. 2003. Accounting for parameter uncertainty in simulation input modeling. *IIE Transactions* **35**(9) 781–792.

Zouaoui, Faker, James R Wilson. 2004. Accounting for input-model and input-parameter uncertainties in simulation. *IIE Transactions* **36**(11) 1135–1151.

Appendix

EC.1. Technical Proofs

Proof of Theorem 1. Let $\delta(y)$ be the delta measure at y . For each $i = 1, \dots, m$, define

$$\tilde{P}^i = \sum_{j=1}^{n^i} \frac{L^i(y_j^i)}{\sum_{r=1}^{n^i} L^i(y_r^i)} \delta(y_j^i)$$

i.e., the distribution with point mass $L^i(y_j^i)/\sum_{r=1}^{n^i} L^i(y_r^i)$ on each y_j^i , where $L^i = dP_0^i/dQ^i$. We first show that as $n \rightarrow \infty$, the solution $(\tilde{P}^i)_{i=1, \dots, m}$ is feasible for the optimization problems in (8) in an appropriate sense.

Consider Case 1. For each $l = 1, \dots, s^i$, by a change measure we have $E_{Q^i}[f_l^i(X^i)L(X^i)] = E_{P_0^i}[f_l^i(X^i)] < \infty$ by our assumption. Also note that $E_{Q^i}L^i = 1$. Therefore, by the law of large numbers,

$$E_{\tilde{P}^i}[f_l^i(X^i)] = \frac{\sum_{j=1}^{n^i} L^i(y_j^i)f_l^i(y_j^i)}{\sum_{j=1}^{n^i} L^i(y_j^i)} = \frac{(1/n^i)\sum_{j=1}^{n^i} L^i(y_j^i)f_l^i(y_j^i)}{(1/n^i)\sum_{j=1}^{n^i} L^i(y_j^i)} \rightarrow E_{Q^i}[f_l^i(X^i)L(X^i)] \text{ a.s.}$$

Since $E_{Q^i}[f_l^i(X^i)L(X^i)] = E_{P_0^i}[f_l^i(X^i)] < \mu_l^i$ by our assumption, we have $E_{\tilde{P}^i}[f_l^i(X^i)] \leq \mu_l^i$ eventually as $n^i \rightarrow \infty$.

Consider Case 2. We have

$$\begin{aligned} d_\phi(\tilde{P}^i, \hat{P}_b^i) &= \sum_{j=1}^{n^i} \phi \left(\frac{L^i(y_j^i)/\sum_{r=1}^{n^i} L^i(y_r^i)}{L_b^i(y_j^i)/\sum_{r=1}^{n^i} L_b^i(y_r^i)} \right) \frac{L_b^i(y_j^i)}{\sum_{r=1}^{n^i} L_b^i(y_r^i)} \\ &= \frac{1}{n^i} \sum_{j=1}^{n^i} \phi \left(\tilde{L}^i(y_j^i) \frac{(1/n^i)\sum_{r=1}^{n^i} L_b^i(y_r^i)}{(1/n^i)\sum_{r=1}^{n^i} L^i(y_r^i)} \right) \frac{L_b^i(y_j^i)}{(1/n^i)\sum_{r=1}^{n^i} L_b^i(y_r^i)} \end{aligned}$$

where $\tilde{L}^i = dP_0^i/dP_b^i$. Consider, for a given $\epsilon > 0$,

$$\begin{aligned} &P(|d_\phi(\tilde{P}^i, \hat{P}_b^i) - d_\phi(P_0^i, P_b^i)| > \epsilon) \\ &= P \left(\left| \frac{1}{n^i} \sum_{j=1}^{n^i} \phi \left(\tilde{L}^i(y_j^i) \frac{(1/n^i)\sum_{r=1}^{n^i} L_b^i(y_r^i)}{(1/n^i)\sum_{r=1}^{n^i} L^i(y_r^i)} \right) \frac{L_b^i(y_j^i)}{(1/n^i)\sum_{r=1}^{n^i} L_b^i(y_r^i)} - d_\phi(P_0^i, P_b^i) \right| > \epsilon \right) \\ &\leq P \left(\left| \frac{1}{n^i} \sum_{j=1}^{n^i} \phi \left(\tilde{L}^i(y_j^i) \frac{(1/n^i)\sum_{r=1}^{n^i} L_b^i(y_r^i)}{(1/n^i)\sum_{r=1}^{n^i} L^i(y_r^i)} \right) \frac{L_b^i(y_j^i)}{(1/n^i)\sum_{r=1}^{n^i} L_b^i(y_r^i)} - \frac{1}{n^i} \sum_{j=1}^{n^i} \phi(\tilde{L}^i(y_j^i)) \frac{L_b^i(y_j^i)}{(1/n^i)\sum_{r=1}^{n^i} L_b^i(y_r^i)} \right| > \frac{\epsilon}{2} \right) \\ &\quad + P \left(\left| \frac{1}{n^i} \sum_{j=1}^{n^i} \phi(\tilde{L}^i(y_j^i)) \frac{L_b^i(y_j^i)}{(1/n^i)\sum_{r=1}^{n^i} L_b^i(y_r^i)} - d_\phi(P_0^i, P_b^i) \right| > \frac{\epsilon}{2} \right) \end{aligned} \tag{EC.1}$$

We analyze the two terms in (EC.1). For any sufficiently small $\lambda > 0$, the first term is bounded from above by

$$\begin{aligned}
& P \left(\left| \frac{1}{n^i} \sum_{j=1}^{n^i} \phi \left(\tilde{L}^i(y_j^i) \frac{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)}{(1/n^i) \sum_{r=1}^{n^i} L^i(y_r^i)} \right) \frac{L_b^i(y_j^i)}{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)} - \frac{1}{n^i} \sum_{j=1}^{n^i} \phi(\tilde{L}^i(y_j^i)) \frac{L_b^i(y_j^i)}{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)} \right| > \frac{\epsilon}{2} \right. \\
& \left. ; \left| \frac{1}{n^i} \sum_{r=1}^{n^i} L^i(y_r^i) - 1 \right| \leq \lambda, \left| \frac{1}{n^i} \sum_{r=1}^{n^i} L_b^i(y_r^i) - 1 \right| \leq \lambda \right) \\
& + P \left(\left| \frac{1}{n^i} \sum_{j=1}^{n^i} \phi \left(\tilde{L}^i(y_j^i) \frac{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)}{(1/n^i) \sum_{r=1}^{n^i} L^i(y_r^i)} \right) \frac{L_b^i(y_j^i)}{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)} - \frac{1}{n^i} \sum_{j=1}^{n^i} \phi(\tilde{L}^i(y_j^i)) \frac{L_b^i(y_j^i)}{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)} \right| > \frac{\epsilon}{2} \right. \\
& \left. ; \left| \frac{1}{n^i} \sum_{r=1}^{n^i} L^i(y_r^i) - 1 \right| > \lambda \text{ or } \left| \frac{1}{n^i} \sum_{r=1}^{n^i} L_b^i(y_r^i) - 1 \right| > \lambda \right) \\
& \leq P \left(\frac{1}{n^i} \sum_{j=1}^{n^i} (|\phi(\tilde{L}^i(y_j^i))| + 1) O(\lambda) \frac{L_b^i(y_j^i)}{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)} > \frac{\epsilon}{2}; \left| \frac{1}{n^i} \sum_{r=1}^{n^i} L^i(y_r^i) - 1 \right| \leq \lambda, \left| \frac{1}{n^i} \sum_{r=1}^{n^i} L_b^i(y_r^i) - 1 \right| \leq \lambda \right) \\
& + P \left(\left| \frac{1}{n^i} \sum_{r=1}^{n^i} L^i(y_r^i) - 1 \right| > \lambda \text{ or } \left| \frac{1}{n^i} \sum_{r=1}^{n^i} L_b^i(y_r^i) - 1 \right| > \lambda \right) \tag{EC.2}
\end{aligned}$$

where the first term in the last inequality follows from the continuity condition on ϕ , with $O(\lambda)$ being a deterministic positive function of λ that converges to 0 as $\lambda \rightarrow 0$. This first term is further bounded from above by

$$P \left(\frac{1}{n^i} \sum_{j=1}^{n^i} (|\phi(\tilde{L}^i(y_j^i))| + 1) \frac{L_b^i(y_j^i)}{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)} O(\lambda) > \frac{\epsilon}{2} \right) \tag{EC.3}$$

By the law of large numbers, we have

$$\frac{1}{n^i} \sum_{j=1}^{n^i} (|\phi(\tilde{L}^i(y_j^i))| + 1) L_b^i(y_j^i) \rightarrow E_{Q^i}[(|\phi(\tilde{L}^i(X^i))| + 1) L_b^i(X^i)] = E_{P_b^i}[\phi(\tilde{L}^i(X^i))| + 1] \text{ a.s.}$$

by using our assumption $E_{P_b^i}[\phi(\tilde{L}^i(X^i))| < \infty$. Moreover, by the law of large numbers again, we have $(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i) \rightarrow 1$ a.s.. Thus,

$$\frac{1}{n^i} \sum_{j=1}^{n^i} (|\phi(\tilde{L}^i(y_j^i))| + 1) \frac{L_b^i(y_j^i)}{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)} \rightarrow E_{P_b^i}[\phi(\tilde{L}^i(X^i))| + 1 \text{ a.s.}]$$

When λ is chosen small enough relative to $\epsilon/2$, we have (EC.3) go to 0 as $n^i \rightarrow \infty$.

Since both $\frac{1}{n^i} \sum_{r=1}^{n^i} L^i(y_r^i) \rightarrow 1$ and $\frac{1}{n^i} \sum_{r=1}^{n^i} L_b^i(y_r^i) \rightarrow 1$ a.s., the second term in (EC.2) also goes to 0 as $n^i \rightarrow \infty$. This concludes that the first term in (EC.1) goes to 0.

For the second term in (EC.1), note that

$$\frac{1}{n^i} \sum_{j=1}^{n^i} \phi(\tilde{L}^i(y_j^i)) L_b^i(y_j^i) \rightarrow E_{Q^i}[\phi(\tilde{L}^i(X^i)) L_b^i(X^i)] = E_{P_b^i}[\phi(\tilde{L}^i(X^i))] = d_\phi(P_0^i, P_b^i) \text{ a.s.}$$

by the law of large numbers and the assumption that $E_{P_b^i}|\phi(\tilde{L}^i(X^i))| < \infty$. Moreover, since $(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i) \rightarrow 1$, we get

$$\frac{1}{n^i} \sum_{j=1}^{n^i} \phi(\tilde{L}^i(y_j^i)) \frac{L_b^i(y_j^i)}{(1/n^i) \sum_{r=1}^{n^i} L_b^i(y_r^i)} \rightarrow d_\phi(P_0^i, P_b^i) \text{ a.s.}$$

Thus, the second term in (EC.1) goes to 0 as $n^i \rightarrow \infty$. Therefore, we conclude that $d_\phi(\tilde{P}^i, \hat{P}_b^i) \xrightarrow{p} d_\phi(P_0^i, P_b^i)$. Since $d_\phi(P_0^i, P_b^i) < \eta^i$ by our assumption, we have $P(d_\phi(\tilde{P}^i, \hat{P}_b^i) \leq \eta^i) \rightarrow 1$ as $n^i \rightarrow \infty$.

Next we consider the objective in (8). We show that $Z(\tilde{P}^1, \dots, \tilde{P}^m) - Z(P_0^1, \dots, P_0^m) = O_p(1/\sqrt{n})$, following the argument in the theory of differentiable statistical functionals (e.g., Serfling (2009), Chapter 6). For any λ between 0 and 1, we write

$$\begin{aligned} & Z(P_0^1 + \lambda(\tilde{P}^1 - P_0^1), \dots, P_0^m + \lambda(\tilde{P}^m - P_0^m)) \\ &= \int \dots \int h(\mathbf{x}^1, \dots, \mathbf{x}^m) \prod_{i=1}^m \prod_{t=1}^{T^i} d[P_0^i + \lambda(\tilde{P}^i - P_0^i)](x_t^i) \\ &= \sum_{k=0}^T \lambda^k \sum_{u \in \mathcal{I}^k} \int \dots \int h(\mathbf{x}^1, \dots, \mathbf{x}^m) \prod_{(i,t) \in (\mathcal{S}_u^k)^c} dP_0^i(x_t^i) \prod_{(i,t) \in \mathcal{S}_u^k} d(\tilde{P}^i - P_0^i)(x_t^i) \end{aligned}$$

where $\{\mathcal{S}_u^k\}_{u \in \mathcal{I}^k}$ is the collection of all subsets of $\{(i, t) : i = 1, \dots, m, t = 1, \dots, T^i\}$ with cardinality k , and \mathcal{I}^k indexes all these subsets. Note that

$$\begin{aligned} & \frac{d}{d\lambda} Z(P_0^1 + \lambda(\tilde{P}^1 - P_0^1), \dots, P_0^m + \lambda(\tilde{P}^m - P_0^m)) \Big|_{\lambda=0^+} \\ &= \sum_{i=1}^m \sum_{t=1}^{T^i} \int \dots \int h(\mathbf{x}^1, \dots, \mathbf{x}^m) \prod_{(j,s):(j,s) \neq (i,t)} dP_0^j(x_s^j) d(\tilde{P}^i - P_0^i)(x_t^i) \\ &= \sum_{i=1}^m \int \varphi^i(x; P_0^1, \dots, P_0^m) d(\tilde{P}^i - P_0^i)(x) \end{aligned} \tag{EC.4}$$

where

$$\varphi^i(x; P_0^1, \dots, P_0^m) = \sum_{t=1}^{T^i} E_{P_0^1, \dots, P_0^m}[h(\mathbf{X}^1, \dots, \mathbf{X}^m) | X_t^i = x] \tag{EC.5}$$

By the definition of L^i , we can write (EC.4) as

$$\begin{aligned} & \sum_{i=1}^m \left(\frac{\int \varphi^i(x; P_0^1, \dots, P_0^m) L^i(x) d\hat{Q}^i(x)}{(1/n^i) \sum_{j=1}^{n^i} L^i(y_j^i)} - \int \varphi^i(x; P_0^1, \dots, P_0^m) L^i(x) dQ^i(x) \right) \\ &= \sum_{i=1}^m \left(\frac{\int \varphi^i(x; P_0^1, \dots, P_0^m) L^i(x) d(\hat{Q}^i - Q^i)(x)}{(1/n^i) \sum_{j=1}^{n^i} L^i(y_j^i)} - \int \varphi^i(x; P_0^1, \dots, P_0^m) L^i(x) dQ^i(x) \left(1 - \frac{1}{(1/n^i) \sum_{j=1}^{n^i} L^i(y_j^i)} \right) \right) \end{aligned}$$

where \hat{Q}^i is the empirical distribution $(1/n^i) \sum_{j=1}^{n^i} \delta(y_j^i)$ on the n^i observations generated from Q^i .

Suppose $\varphi^i(x; P_0^1, \dots, P_0^m) L^i(x) = 0$ a.s., then $\int \varphi^i(x; P_0^1, \dots, P_0^m) L^i(x) d(\hat{Q}^i - Q^i)(x) = 0$ a.s..

Otherwise, using the assumed boundedness of h , hence $\varphi^i(x; P_0^1, \dots, P_0^m)$, and L^i , we have, by the central limit theorem,

$$\sqrt{n^i} \left(\int \varphi^i(x; P_0^1, \dots, P_0^m) L^i(x) d(\hat{Q}^i - Q^i)(x) \right) \Rightarrow N(0, (\sigma^i)^2)$$

where $(\sigma^i)^2 = \text{Var}_{Q^i}(\varphi^i(X^i; P_0^1, \dots, P_0^m) L^i(X^i)) > 0$ is finite. Since $(1/n^i) \sum_{j=1}^{n^i} L^i(y_j^i) \rightarrow 1$ a.s. by the law of large numbers, and that $\int \varphi^i(x; P_0^1, \dots, P_0^m) L^i(x) d\hat{Q}^i(x)$ is bounded, the second term in (EC.6) converges to 0 a.s.. Thus, by Slutsky's theorem, each summand in (EC.6) converges in distribution to $N(0, (\sigma^i)^2)$. Since for each i we have $n^i = nw^i$ for some fixed $w^i > 0$, we conclude that (EC.6) equal $O_p(1/\sqrt{n})$.

Now consider

$$\begin{aligned} & \frac{d^2}{d\lambda^2} Z(P_0^1 + \lambda(\tilde{P}^1 - P_0^1), \dots, P_0^m + \lambda(\tilde{P}^m - P_0^m)) \\ &= \sum_{k=2}^T k(k-1)\lambda^{k-2} \sum_{u \in \mathcal{I}^k} \int \dots \int h(\mathbf{x}^1, \dots, \mathbf{x}^m) \prod_{(i,t) \in (\mathcal{S}_u^k)^c} dP_0^i(x_t^i) \prod_{(i,t) \in \mathcal{S}_u^k} d(\tilde{P}^i - P_0^i)(x_t^i) \end{aligned} \quad (\text{EC.7})$$

Fixing each \mathcal{S}_u^k , we define

$$h_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) = \int \dots \int h(\mathbf{x}^1, \dots, \mathbf{x}^m) \prod_{(i,t) \in (\mathcal{S}_u^k)^c} dP_0^i(x_t^i)$$

where $\mathbf{x}_{\mathcal{S}_u^k} = (x_t^i)_{(i,t) \in \mathcal{S}_u^k}$. Next define

$$\begin{aligned} & \tilde{h}_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) \\ &= h_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) - \sum_{(j,t) \in \mathcal{S}_u^k} \int h_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) dP_0^j(x_t^j) + \sum_{(j_1,t_1), (j_2,t_2) \in \mathcal{S}_u^k} \int \int h_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) dP_0^{j_1}(x_{t_1}^{j_1}) dP_0^{j_2}(x_{t_2}^{j_2}) - \dots \\ &+ (-1)^k \int \dots \int h_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) dP_0^{j_1}(x_{t_1}^{j_1}) \dots dP_0^{j_k}(x_{t_k}^{j_k}) \end{aligned}$$

where each summation above is over the set of all possible combinations of $(j, t) \in \mathcal{S}_u^k$ with increasing size. Direct verification shows that $\tilde{h}_{\mathcal{S}_u^k}$ has the property that

$$\int \cdots \int \tilde{h}_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) \prod_{(i,t) \in \mathcal{S}_u^k} dR^j(x_t^j) = \int \cdots \int h_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) \prod_{(i,t) \in \mathcal{S}_u^k} d(R^j(x_t^j) - P_0^j(x_t^j))$$

for any probability measures R^j 's, and

$$\int \tilde{h}_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) dP_0^j(x_t^j) = 0 \quad (\text{EC.8})$$

for any $(j, t) \in \mathcal{S}_u^k$. Thus, (EC.7) is equal to

$$\sum_{k=2}^T k(k-1)\lambda^{k-2} \sum_{u \in \mathcal{I}^k} \int \tilde{h}_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) \prod_{(i,t) \in \mathcal{S}_u^k} d\tilde{P}^i(x_t^i)$$

Now, viewing \tilde{P}^i as randomly generated from Q^i , consider

$$\begin{aligned} & E_{Q^1, \dots, Q^m} \left(\sum_{k=2}^T k(k-1)\lambda^{k-2} \sum_{u \in \mathcal{I}^k} \int \tilde{h}_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) \prod_{(i,t) \in \mathcal{S}_u^k} d\tilde{P}^i(x_t^i) \right)^2 \\ &= E_{Q^1, \dots, Q^m} \left[\left(\sum_{k=2}^T k(k-1)\lambda^{k-2} \sum_{u \in \mathcal{I}^k} \int \tilde{h}_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) \prod_{(i,t) \in \mathcal{S}_u^k} d\tilde{P}^i(x_t^i) \right)^2 ; \frac{1}{n^i} \sum_{r=1}^{n^i} L^i(Y_r^i) \geq 1 - \epsilon \text{ for all } i = 1, \dots, m \right] \\ &+ E_{Q^1, \dots, Q^m} \left[\left(\sum_{k=2}^T k(k-1)\lambda^{k-2} \sum_{u \in \mathcal{I}^k} \int \tilde{h}_{\mathcal{S}_u^k}(\mathbf{x}_{\mathcal{S}_u^k}) \prod_{(i,t) \in \mathcal{S}_u^k} d\tilde{P}^i(x_t^i) \right)^2 ; \frac{1}{n^i} \sum_{r=1}^{n^i} L^i(Y_r^i) < 1 - \epsilon \text{ for some } i = 1, \dots, m \right] \end{aligned} \quad (\text{EC.9})$$

We analyze the two terms in (EC.9). Note that the first term can be written as

$$\begin{aligned} & E_{Q^1, \dots, Q^m} \left[\left(\sum_{k=2}^T k(k-1)\lambda^{k-2} \sum_{u \in \mathcal{I}^k} \frac{1}{n^{i_1} n^{i_2} \dots n^{i_k}} \sum_{j_1=1}^{n^{i_1}} \dots \sum_{j_k=1}^{n^{i_k}} \tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) \frac{L^{i_1}(Y_{j_1}^{i_1}) \dots L^{i_k}(Y_{j_k}^{i_k})}{\prod_{s=1}^k ((1/n^{i_s}) \sum_{r=1}^{n^{i_s}} L^{i_s}(Y_r^{i_s}))} \right)^2 ; \right. \\ & \quad \left. \frac{1}{n^i} \sum_{r=1}^{n^i} L^i(Y_r^i) \geq 1 - \epsilon \text{ for all } i = 1, \dots, m \right] \\ & \leq \left(\sum_{k=2}^T k(k-1)\lambda^{k-2} \sum_{u \in \mathcal{I}^k} \frac{1}{n^{i_1} n^{i_2} \dots n^{i_k}} \left(E_{Q^1, \dots, Q^m} \left[\left(\sum_{j_1=1}^{n^{i_1}} \dots \sum_{j_k=1}^{n^{i_k}} \tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) \frac{L^{i_1}(Y_{j_1}^{i_1}) \dots L^{i_k}(Y_{j_k}^{i_k})}{\prod_{s=1}^k ((1/n^{i_s}) \sum_{r=1}^{n^{i_s}} L^{i_s}(Y_r^{i_s}))} \right)^2 ; \right. \right. \\ & \quad \left. \left. \frac{1}{n^i} \sum_{r=1}^{n^i} L^i(Y_r^i) \geq 1 - \epsilon \text{ for all } i = 1, \dots, m \right] \right)^{1/2} \right)^2 \end{aligned} \quad (\text{EC.10})$$

by Minkowski's inequality, where we view Y_j^i 's as the random variables constituting the observations generated from Q^i 's. Since the expression $\prod_{s=1}^k ((1/n^{i_s}) \sum_{r=1}^{n^{i_s}} L^{i_s}(Y_r^{i_s}))$ inside the expectation in (EC.10) does not depend on the j_s 's, (EC.10) is further bounded from above by

$$\left(\sum_{k=2}^T k(k-1)\lambda^{k-2} \sum_{u \in \mathcal{I}^k} \frac{1}{n^{i_1} n^{i_2} \dots n^{i_k}} \left(E_{Q^1, \dots, Q^m} \left(\sum_{j_1=1}^{n^{i_1}} \dots \sum_{j_k=1}^{n^{i_k}} \tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) \frac{L^{i_1}(Y_{j_1}^{i_1}) \dots L^{i_k}(Y_{j_k}^{i_k})}{(1-\epsilon)^k} \right)^2 \right)^{1/2} \right)^2$$

$$= \left(\sum_{k=2}^T \frac{k(k-1)\lambda^{k-2}}{(1-\epsilon)^k} \sum_{u \in \mathcal{I}^k} \frac{1}{n^{i_1} n^{i_2} \cdots n^{i_k}} \left(\sum_{j_1=1}^{n^{i_1}} \cdots \sum_{j_k=1}^{n^{i_k}} \sum_{j'_1=1}^{n^{i_1}} \cdots \sum_{j'_k=1}^{n^{i_k}} E_{Q^1, \dots, Q^m} \left[\tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) L^{i_1}(Y_{j_1}^{i_1}) \cdots L^{i_k}(Y_{j_k}^{i_k}) \right. \right. \right. \\ \left. \left. \left. \tilde{h}_{\mathcal{S}_u^k}(Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k}) L^{i_1}(Y_{j'_1}^{i_1}) \cdots L^{i_k}(Y_{j'_k}^{i_k}) \right] \right] \right)^{1/2} \right)^2 \quad (\text{EC.11})$$

Note that

$$E_{Q^1, \dots, Q^m} \left[\tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) L^{i_1}(Y_{j_1}^{i_1}) \cdots L^{i_k}(Y_{j_k}^{i_k}) \tilde{h}_{\mathcal{S}_u^k}(Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k}) L^{i_1}(Y_{j'_1}^{i_1}) \cdots L^{i_k}(Y_{j'_k}^{i_k}) \right] = 0 \quad (\text{EC.12})$$

if any Y_j^i shows up only once among all those in both $\tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k})$ and $\tilde{h}_{\mathcal{S}_u^k}(Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k})$ in the expectation. To see this, suppose without loss of generality that $Y_{j_1}^{i_1}$ appears only once. Then we have

$$E_{Q^1, \dots, Q^m} \left[\tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) L^{i_1}(Y_{j_1}^{i_1}) \cdots L^{i_k}(Y_{j_k}^{i_k}) \tilde{h}_{\mathcal{S}_u^k}(Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k}) L^{i_1}(Y_{j'_1}^{i_1}) \cdots L^{i_k}(Y_{j'_k}^{i_k}) \right] \\ = E_{Q^1, \dots, Q^m} \left[E_{Q^1, \dots, Q^m} \left[\tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) L^{i_1}(Y_{j_1}^{i_1}) \middle| Y_{j_2}^{i_2}, \dots, Y_{j_k}^{i_k}, Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k} \right] L^{i_2}(Y_{j_2}^{i_2}) \cdots L^{i_k}(Y_{j_k}^{i_k}) \right. \\ \left. \tilde{h}_{\mathcal{S}_u^k}(Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k}) L^{i_1}(Y_{j'_1}^{i_1}) \cdots L^{i_k}(Y_{j'_k}^{i_k}) \right] \\ = E_{Q^1, \dots, Q^m} \left[E_{P_0^{i_1}} \left[\tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) \middle| Y_{j_2}^{i_2}, \dots, Y_{j_k}^{i_k}, Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k} \right] L^{i_2}(Y_{j_2}^{i_2}) \cdots L^{i_k}(Y_{j_k}^{i_k}) \right. \\ \left. \tilde{h}_{\mathcal{S}_u^k}(Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k}) L^{i_1}(Y_{j'_1}^{i_1}) \cdots L^{i_k}(Y_{j'_k}^{i_k}) \right] \\ = 0$$

since $E_{P_0^{i_1}} \left[\tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) \middle| Y_{j_2}^{i_2}, \dots, Y_{j_k}^{i_k}, Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k} \right] = 0$ by (EC.8).

The observation in (EC.12) implies that the summation in (EC.11)

$$\sum_{j_1=1}^{n^{i_1}} \cdots \sum_{j_k=1}^{n^{i_k}} \sum_{j'_1=1}^{n^{i_1}} \cdots \sum_{j'_k=1}^{n^{i_k}} E_{Q^1, \dots, Q^m} \left[\tilde{h}_{\mathcal{S}_u^k}(Y_{j_1}^{i_1}, \dots, Y_{j_k}^{i_k}) L^{i_1}(Y_{j_1}^{i_1}) \cdots L^{i_k}(Y_{j_k}^{i_k}) \tilde{h}_{\mathcal{S}_u^k}(Y_{j'_1}^{i_1}, \dots, Y_{j'_k}^{i_k}) L^{i_1}(Y_{j'_1}^{i_1}) \cdots L^{i_k}(Y_{j'_k}^{i_k}) \right]$$

contains only $O(n^k)$ non-zero summands. This is because in each non-zero summand only at most k distinct Y_j^i 's can be present inside the expectation, and the cardinality of such combinations is $O(n^k)$. Note that each summand is bounded since h , hence $\tilde{h}_{\mathcal{S}_u^k}$, and L^i are all bounded by our assumptions. Hence (EC.11) is

$$\left(\sum_{k=2}^T \frac{k(k-1)\lambda^{k-2}}{(1-\epsilon)^k} \binom{T}{k} O\left(\frac{1}{n^{k/2}}\right) \right)^2 = O\left(\frac{1}{n^2}\right) \quad (\text{EC.13})$$

This shows that (EC.7) is $O_p(1/n)$ for any λ between 0 and 1. Therefore, by using Taylor's expansion, and the conclusion that (EC.6) is $O_p(1/\sqrt{n})$, we have

$$Z(\tilde{P}^1, \dots, \tilde{P}^m) = Z(P_0^1, \dots, P_0^m) + O_p\left(\frac{1}{\sqrt{n}}\right) = Z_0 + O_p\left(\frac{1}{\sqrt{n}}\right) \quad (\text{EC.14})$$

Note that we have shown previously that $P(\tilde{P}^i \in \hat{\mathcal{U}}^i) \rightarrow 1$ for any $i = 1, \dots, m$ in both Cases 1 and 2. Using this and (EC.14), for any given $\epsilon > 0$, we can choose $M, N > 0$ big enough such that

$$P(\sqrt{n}(\hat{Z}_* - Z_0) > M) \leq P(|\sqrt{n}(Z(\tilde{P}^1, \dots, \tilde{P}^i) - Z_0)| > M) + \sum_{i=1}^m P(\tilde{P}^i \notin \hat{\mathcal{U}}^i) < \epsilon$$

and similarly

$$P(\sqrt{n}(Z_0 - \hat{Z}^*) > M) \leq P(|\sqrt{n}(Z(\tilde{P}^1, \dots, \tilde{P}^i) - Z_0)| > M) + \sum_{i=1}^m P(\tilde{P}^i \notin \hat{\mathcal{U}}^i) < \epsilon$$

for any $n > N$. This concludes that

$$\hat{Z}_* \leq Z_0 + O_p\left(\frac{1}{\sqrt{n}}\right) \leq \hat{Z}^*$$

Proof of Theorem 3. To prove 1., consider first a mixture of $\mathbf{p}^i = (p_j^i)_{j=1, \dots, n^i}$ with an arbitrary $\mathbf{q}^i \in \mathcal{P}_{n^i}$, in the form $(1 - \epsilon)\mathbf{p}^i + \epsilon\mathbf{q}^i$. It satisfies

$$\frac{d}{d\epsilon} Z(\mathbf{p}^1, \dots, \mathbf{p}^{i-1}, (1 - \epsilon)\mathbf{p}^i + \epsilon\mathbf{q}^i, \mathbf{p}^{i+1}, \dots, \mathbf{p}^m) \Big|_{\epsilon=0} = \nabla^i Z(\mathbf{p})'(\mathbf{q}^i - \mathbf{p}^i)$$

by the chain rule. In particular, we must have

$$\psi_j^i(\mathbf{p}) = \nabla^i Z(\mathbf{p})'(\mathbf{1}_j^i - \mathbf{p}^i) = \partial_j^i Z(\mathbf{p}) - \nabla^i Z(\mathbf{p})' \mathbf{p}^i \quad (\text{EC.15})$$

where $\partial_j^i Z(\mathbf{p})$ denotes partial derivative of Z with respect to p_j^i . Writing (EC.15) for all j together gives

$$\psi^i(\mathbf{p}) = \nabla^i Z(\mathbf{p}) - (\nabla^i Z(\mathbf{p})' \mathbf{p}^i) \mathbf{1}^i$$

where $\mathbf{1}^i \in \mathbb{R}^{n^i}$ is a vector of 1. Therefore

$$\psi^i(\mathbf{p})'(\mathbf{q}^i - \mathbf{p}^i) = (\nabla^i Z(\mathbf{p}) - (\nabla^i Z(\mathbf{p})' \mathbf{p}^i) \mathbf{1}^i)'(\mathbf{q}^i - \mathbf{p}^i) = \nabla^i Z(\mathbf{p})'(\mathbf{q}^i - \mathbf{p}^i)$$

since $\mathbf{q}^i, \mathbf{p}^i \in \mathcal{P}_{n^i}$. Summing up over i , (12) follows.

To prove 2., note that we have

$$\begin{aligned}\psi_j^i(\mathbf{p}) &= \frac{d}{d\epsilon} Z(\mathbf{p}^1, \dots, \mathbf{p}^{i-1}, (1-\epsilon)\mathbf{p}^i + \epsilon\mathbf{1}_j^i, \mathbf{p}^{i+1}, \dots, \mathbf{p}^m) \Big|_{\epsilon=0} \\ &= \frac{d}{d\epsilon} E_{\mathbf{p}^1, \dots, \mathbf{p}^{i-1}, (1-\epsilon)\mathbf{p}^i + \epsilon\mathbf{1}_j^i, \mathbf{p}^{i+1}, \dots, \mathbf{p}^m} [h(\mathbf{X})] \Big|_{\epsilon=0} \\ &= E_{\mathbf{p}} [h(\mathbf{X}) s_j^i(\mathbf{X}^i)]\end{aligned}\tag{EC.16}$$

where $s_j^i(\cdot)$ is the score function defined as

$$s_j^i(\mathbf{x}^i) = \sum_{t=1}^{T^i} \frac{d}{d\epsilon} \log((1-\epsilon)p^i(x_t^i) + \epsilon I(x_t^i = y_j^i)) \Big|_{\epsilon=0}.\tag{EC.17}$$

Here $p^i(x_t^i) = p_j^i$ where j is chosen such that $x_t^i = y_j^i$. The last equality in (EC.16) follows from the fact that

$$\frac{d}{d\epsilon} \prod_{t=1}^{T^i} ((1-\epsilon)p^i(x_t^i) + \epsilon I(x_t^i = y_j^i)) \Big|_{\epsilon=0} = \frac{d}{d\epsilon} \sum_{t=1}^{T^i} \log((1-\epsilon)p^i(x_t^i) + \epsilon I(x_t^i = y_j^i)) \Big|_{\epsilon=0} \cdot \prod_{t=1}^{T^i} p^i(x_t^i)$$

Note that (EC.17) can be further written as

$$\sum_{t=1}^{T^i} \frac{-p^i(x_t^i) + I(x_t^i = y_j^i)}{p^i(x_t^i)} = -T^i + \sum_{t=1}^{T^i} \frac{I(x_t^i = y_j^i)}{p^i(x_t^i)} = -T^i + \sum_{t=1}^{T^i} \frac{I(x_t^i = y_j^i)}{p_j^i}$$

which leads to (13).

Proof of Lemma 1 We have

$$\begin{aligned}Var_{\mathbf{p}}(h(\mathbf{X}) s_j^i(\mathbf{X}^i)) &\leq E_{\mathbf{p}}(h(\mathbf{X}) s_j^i(\mathbf{X}^i))^2 \leq M^2 E_{\mathbf{p}}(s_j^i(\mathbf{X}^i))^2 = M^2 (Var_{\mathbf{p}}(s_j^i(\mathbf{X}^i)) + (E_{\mathbf{p}}[s_j^i(\mathbf{X}^i)])^2)\end{aligned}\tag{EC.18}$$

Now note that by the definition of $s_j^i(\mathbf{X})$ in (14) we have $E_{\mathbf{p}}[s_j^i(\mathbf{X}^i)] = 0$ and

$$Var_{\mathbf{p}}(s_j^i(\mathbf{X}^i)) = \frac{T^i Var_{\mathbf{p}}(I(X_t^i = y_j^i))}{(p_j^i)^2} = \frac{T^i (1 - p_j^i)}{p_j^i}$$

Hence, from (EC.18), we conclude that $Var_{\mathbf{p}}(h(\mathbf{X}) s_j^i(\mathbf{X}^i)) \leq M^2 T^i (1 - p_j^i) / p_j^i$.

Proof of Proposition 3 Consider the Lagrangian relaxation

$$\begin{aligned}
& \max_{\alpha \geq 0, \lambda \in \mathbb{R}} \min_{\mathbf{p}^i \geq \mathbf{0}} \sum_{j=1}^{n^i} p_j^i \xi_j + \alpha \left(\sum_{j=1}^{n^i} p_{b,j}^i \phi \left(\frac{p_j^i}{p_{b,j}^i} \right) - \eta^i \right) + \lambda \left(\sum_{j=1}^{n^i} p_j^i - 1 \right) \quad (\text{EC.19}) \\
&= \max_{\alpha \geq 0, \lambda \in \mathbb{R}} -\alpha \sum_{j=1}^{n^i} p_{b,j}^i \max_{p_j^i \geq 0} \left\{ -\frac{\xi_j + \lambda}{\alpha} \frac{p_j^i}{p_{b,j}^i} - \phi \left(\frac{p_j^i}{p_{b,j}^i} \right) \right\} - \alpha \eta^i - \lambda \\
&= \max_{\alpha \geq 0, \lambda \in \mathbb{R}} -\alpha \sum_{j=1}^{n^i} p_{b,j}^i \phi^* \left(-\frac{\xi_j + \lambda}{\alpha} \right) - \alpha \eta^i - \lambda
\end{aligned}$$

In the particular case that $\alpha^* = 0$, the optimal value of (EC.19) is the same as

$$\max_{\lambda \in \mathbb{R}} \min_{\mathbf{p}^i \geq \mathbf{0}} \sum_{j=1}^{n^i} p_j^i \xi_j + \lambda \left(\sum_{j=1}^{n^i} p_j^i - 1 \right)$$

whose inner minimization is equivalent to $\min_{\mathbf{p}^i \in \mathcal{P}^i} \sum_{j=1}^{n^i} p_j^i \xi_j = \min_{j \in \{1, \dots, n^i\}} \xi_j$. Among all solutions that lead to this objective value, we find the one that solves

$$\min_{p_j^i, j \in \mathcal{M}^i: \sum_{j \in \mathcal{M}^i} p_j^i = 1} \sum_{j \in \mathcal{M}^i} p_{b,j}^i \phi \left(\frac{p_j^i}{p_{b,j}^i} \right) \quad (\text{EC.20})$$

Now note that by the convexity of ϕ and Jensen's inequality, for any $\sum_{j \in \mathcal{M}^i} p_j^i = 1$, we have

$$\sum_{j \in \mathcal{M}^i} p_{b,j}^i \phi \left(\frac{p_j^i}{p_{b,j}^i} \right) = \sum_{r \in \mathcal{M}^i} p_{b,r}^i \sum_{j \in \mathcal{M}^i} \frac{p_{b,j}^i}{\sum_{r \in \mathcal{M}^i} p_{b,r}^i} \phi \left(\frac{p_j^i}{p_{b,j}^i} \right) \geq \sum_{j \in \mathcal{M}^i} p_{b,j}^i \phi \left(\frac{1}{\sum_{j \in \mathcal{M}^i} p_{b,j}^i} \right) = \phi \left(\frac{1}{\sum_{j \in \mathcal{M}^i} p_{b,j}^i} \right) \quad (\text{EC.21})$$

It is easy to see that choosing p_j^i in (EC.20) as q_j^i depicted in (22) achieves the lower bound in (EC.21), hence concluding the proposition.

Proof of Proposition 4 Consider the Lagrangian for the optimization (18)

$$\min_{\mathbf{p}^i \in \mathcal{P}^i} \sum_{j=1}^{n^i} \xi_j p_j^i + \alpha \left(\sum_{j=1}^{n^i} p_j^i \log \frac{p_j^i}{p_{b,j}^i} - \eta^i \right) \quad (\text{EC.22})$$

By Theorem 1, P.220 in Luenberger (1969), suppose that one can find $\alpha^* \geq 0$ such that $\mathbf{q}^i = (q_j^i)_{j=1, \dots, n^i} \in \mathcal{P}_{n^i}$ minimizes (EC.22) for $\alpha = \alpha^*$ and moreover that $\alpha^* \left(\sum_{j=1}^{n^i} q_j^i \log \frac{q_j^i}{p_{b,j}^i} - \eta^i \right) = 0$, then \mathbf{q}^i is optimal for (18).

Suppose $\alpha^* = 0$, then the minimizer of (EC.22) can be any probability distributions that have masses concentrated on the set of indices in \mathcal{M}^i . Any one of these distributions that lies in $\hat{\mathcal{U}}^i$

will be an optimal solution to (18). To check whether any of them lies in $\hat{\mathcal{U}}^i$, consider the one that has the minimum $d_\phi(\mathbf{q}^i, \mathbf{p}_b^i)$ and see whether it is less than or equal to η^i . In other words, we want to find $\min_{p_j^i, j \in \mathcal{M}^i: \sum_{j \in \mathcal{M}^i} p_j^i = 1} \sum_{j \in \mathcal{M}^i} p_j^i \log(p_j^i / p_{b,j}^i)$. The optimal solution to this minimization is $p_{b,j}^i / \sum_{j \in \mathcal{M}^i} p_{b,j}^i$ for $j \in \mathcal{M}^i$, which gives an optimal value $-\log \sum_{j \in \mathcal{M}^i} p_{b,j}^i$. Thus, if $-\log \sum_{j \in \mathcal{M}^i} p_{b,j}^i \leq \eta^i$, we find an optimal solution \mathbf{q}^i to (18) given by (23).

In the case that $\alpha^* = 0$ does not lead to an optimal solution, or equivalently $-\log \sum_{j \in \mathcal{M}^i} p_{b,j}^i > \eta^i$, we consider $\alpha^* > 0$. We write the objective value of (EC.22) with $\alpha = \alpha^*$ as

$$\sum_{j=1}^{n^i} \xi_j p_j^i + \alpha^* \sum_{j=1}^{n^i} p_j^i \log \frac{p_j^i}{p_{b,j}^i} - \alpha^* \eta^i \quad (\text{EC.23})$$

By Jensen's inequality,

$$\sum_{j=1}^{n^i} p_j^i e^{-\xi_j / \alpha^* - \log(p_j^i / p_{b,j}^i)} \geq e^{-\sum_{j=1}^{n^i} \xi_j p_j^i / \alpha^* - \sum_{j=1}^{n^i} p_j^i \log(p_j^i / p_{b,j}^i)}$$

giving

$$\sum_{j=1}^{n^i} \xi_j p_j^i + \alpha^* \sum_{j=1}^{n^i} p_j^i \log \frac{p_j^i}{p_{b,j}^i} \geq -\alpha^* \log \sum_{j=1}^{n^i} p_{b,j}^i e^{-\xi_j / \alpha^*} \quad (\text{EC.24})$$

It is easy to verify that putting p_j^i as

$$q_j^i = \frac{p_{b,j}^i e^{-\xi_j / \alpha^*}}{\sum_{r=1}^{n^i} p_{b,r}^i e^{-\xi_r / \alpha^*}}$$

gives the lower bound in (EC.24). Thus q_j^i minimizes (EC.23). Moreover, $\alpha^* > 0$ can be chosen such that

$$\sum_{j=1}^{n^i} q_j^i \log \frac{q_j^i}{p_{b,j}^i} = -\frac{\sum_{j=1}^{n^i} \xi_j p_{b,j}^i e^{-\xi_j / \alpha^*}}{\alpha^* \sum_{j=1}^{n^i} p_{b,j}^i e^{-\xi_j / \alpha^*}} - \log \sum_{j=1}^{n^i} p_{b,j}^i e^{-\xi_j / \alpha^*} = \eta^i$$

Letting $\beta = -1/\alpha^*$, we obtain (24) and (25). Note that (25) must bear a negative root because of the following. Note that the left hand side of (25) is continuous, and goes to 0 when $\beta \rightarrow 0$. Defining $\xi_* = \min\{\xi_j : j = 1, \dots, n^i\}$, we have, as $\beta \rightarrow -\infty$, $\varphi_{\xi}^i(\beta) = \log \sum_{j=1}^{n^i} p_{b,j}^i e^{\beta \xi_j} = \log \left(\sum_{j \in \mathcal{M}^i} p_{b,j}^i e^{\beta \xi_*} (1 + \sum_{j \notin \mathcal{M}^i} p_{b,j}^i e^{\beta(\xi_j - \xi_*)} / \sum_{j \in \mathcal{M}^i} p_{b,j}^i) \right) = \beta \xi_* + \log \sum_{j \in \mathcal{M}^i} p_{b,j}^i + O(e^{c_1 \beta})$ for some positive constant c_1 , and $\varphi_{\xi}^i'(\beta) = \sum_{j=1}^{n^i} \xi_j p_{b,j}^i e^{\beta \xi_j} / \sum_{j=1}^{n^i} p_{b,j}^i e^{\beta \xi_j} = \xi_* (1 + \sum_{j \notin \mathcal{M}^i} \xi_j p_{b,j}^i e^{\beta(\xi_j - \xi_*)} / \sum_{j \in \mathcal{M}^i} p_{b,j}^i) / (1 + \sum_{j \notin \mathcal{M}^i} p_{b,j}^i e^{\beta(\xi_j - \xi_*)} / \sum_{j \in \mathcal{M}^i} p_{b,j}^i) = \xi_* + O(e^{c_2 \beta})$ for some positive constant c_2 . So $\beta \varphi_{\xi}^i'(\beta) - \varphi_{\xi}^i(\beta) = -\log \sum_{j \in \mathcal{M}^i} p_{b,j}^i + O(e^{(c_1 \wedge c_2) \beta}) > \eta^i$ when β is negative enough.

Proof of Theorem 4 The proof is an adaptation of Blum (1954). Recall that $\mathbf{p}_k = \text{vec}(\mathbf{p}_k^i : i = 1, \dots, m)$ where we write each component of \mathbf{p}_k as $p_{k,j}^i$. Let $N = \sum_{i=1}^m n^i$ be the total counts of support points. Since $h(\mathbf{X})$ is bounded a.s., we have $|h(\mathbf{X})| \leq M$ a.s. for some M . Without loss of generality, we assume that $Z(\mathbf{p}) \geq 0$ for all \mathbf{p} . Also note that $Z(\mathbf{p})$, as a high-dimensional polynomial, is continuous everywhere in $\hat{\mathcal{U}}$.

For notational convenience, we write $\mathbf{d}_k = \mathbf{q}(\mathbf{p}_k) - \mathbf{p}_k$ and $\hat{\mathbf{d}}_k = \hat{\mathbf{q}}(\mathbf{p}_k) - \mathbf{p}_k$, i.e. \mathbf{d}_k is the k -th step best feasible direction given the exact gradient estimate, and $\hat{\mathbf{d}}_k$ is the one with estimated gradient.

Now, given \mathbf{p}_k , consider the iterative update $\mathbf{p}_{k+1} = (1 - \epsilon_k)\mathbf{p}_k + \epsilon_k \hat{\mathbf{q}}(\mathbf{p}_k) = \mathbf{p}_k + \epsilon_k \hat{\mathbf{d}}_k$. We have, by Taylor series expansion,

$$Z(\mathbf{p}_{k+1}) = Z(\mathbf{p}_k) + \epsilon_k \nabla Z(\mathbf{p}_k)' \hat{\mathbf{d}}_k + \frac{\epsilon_k^2}{2} \hat{\mathbf{d}}_k' \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k$$

for some θ_k between 0 and 1. By Theorem 3, we can rewrite the above as

$$Z(\mathbf{p}_{k+1}) = Z(\mathbf{p}_k) + \epsilon_k \psi(\mathbf{p}_k)' \hat{\mathbf{d}}_k + \frac{\epsilon_k^2}{2} \hat{\mathbf{d}}_k' \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k \quad (\text{EC.25})$$

Consider the second term in the right hand side of (EC.25). We can write

$$\begin{aligned} \psi(\mathbf{p}_k)' \hat{\mathbf{d}}_k &= \hat{\psi}(\mathbf{p}_k)' \hat{\mathbf{d}}_k + (\psi(\mathbf{p}_k) - \hat{\psi}(\mathbf{p}_k))' \hat{\mathbf{d}}_k \\ &\leq \hat{\psi}(\mathbf{p}_k)' \hat{\mathbf{d}}_k + (\psi(\mathbf{p}_k) - \hat{\psi}(\mathbf{p}_k))' \hat{\mathbf{d}}_k \quad \text{by the definition of } \hat{\mathbf{d}}_k \\ &= \psi(\mathbf{p}_k)' \mathbf{d}_k + (\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k))' \mathbf{d}_k + (\psi(\mathbf{p}_k) - \hat{\psi}(\mathbf{p}_k))' \hat{\mathbf{d}}_k \\ &= \psi(\mathbf{p}_k)' \mathbf{d}_k + (\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k))' (\mathbf{d}_k - \hat{\mathbf{d}}_k) \end{aligned} \quad (\text{EC.26})$$

Hence (EC.25) and (EC.26) together imply

$$Z(\mathbf{p}_{k+1}) \leq Z(\mathbf{p}_k) + \epsilon_k \psi(\mathbf{p}_k)' \mathbf{d}_k + \epsilon_k (\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k))' (\mathbf{d}_k - \hat{\mathbf{d}}_k) + \frac{\epsilon_k^2}{2} \hat{\mathbf{d}}_k' \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k$$

Let \mathcal{F}_k be the filtration generated by $\mathbf{p}_1, \dots, \mathbf{p}_k$. We then have

$$\begin{aligned} E[Z(\mathbf{p}_{k+1}) | \mathcal{F}_k] &\leq Z(\mathbf{p}_k) + \epsilon_k \psi(\mathbf{p}_k)' \mathbf{d}_k + \epsilon_k E[(\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k))' (\mathbf{d}_k - \hat{\mathbf{d}}_k) | \mathcal{F}_k] \\ &\quad + \frac{\epsilon_k^2}{2} E[\hat{\mathbf{d}}_k' \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k | \mathcal{F}_k] \end{aligned} \quad (\text{EC.27})$$

We analyze (EC.27) term by term. First, since $Z(\mathbf{p})$ is a high-dimensional polynomial and $\hat{\mathcal{U}}$ is a bounded set, the largest eigenvalue of the Hessian matrix $\nabla^2 Z(\mathbf{p})$, for any $\mathbf{p} \in \hat{\mathcal{U}}$, is uniformly bounded by a constant $H > 0$. Hence

$$E[\hat{\mathbf{d}}'_k \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k | \mathcal{F}_k] \leq H E[\|\hat{\mathbf{d}}_k\|^2 | \mathcal{F}_k] \leq V < \infty \quad (\text{EC.28})$$

for some $V > 0$. Now

$$\begin{aligned} & E[(\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k))' (\mathbf{d}_k - \hat{\mathbf{d}}_k) | \mathcal{F}_k] \\ & \leq \sqrt{E[\|\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k)\|^2 | \mathcal{F}_k] E[\|\mathbf{d}_k - \hat{\mathbf{d}}_k\|^2 | \mathcal{F}_k]} \quad \text{by Cauchy-Schwarz inequality} \\ & \leq \sqrt{E[\|\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k)\|^2 | \mathcal{F}_k] E[2(\|\mathbf{d}_k\|^2 + \|\hat{\mathbf{d}}_k\|^2) | \mathcal{F}_k]} \quad \text{by parallelogram law} \\ & \leq \sqrt{8m E[\|\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k)\|^2 | \mathcal{F}_k]} \quad \text{since } \|\mathbf{d}_k\|^2, \|\hat{\mathbf{d}}_k\|^2 \leq 2m \text{ by using the fact that } \mathbf{p}_k, \mathbf{q}(\mathbf{p}_k), \hat{\mathbf{q}}(\mathbf{p}_k) \in \mathcal{P} \\ & \leq \sqrt{\frac{8mM^2T}{R_k} \sum_{i,j} \frac{1 - p_{k,j}^i}{p_{k,j}^i}} \quad \text{by Lemma 1} \\ & \leq M \sqrt{\frac{8mTN}{R_k \min_{i,j} p_{k,j}^i}} \end{aligned} \quad (\text{EC.30})$$

Note that by iterating the update rule $(1 - \epsilon_k) \mathbf{p}_k + \epsilon_k \mathbf{q}_k$, we have

$$\min_{i,j} p_{k,j}^i \geq \prod_{j=1}^{k-1} (1 - \epsilon_j) \delta$$

where $\delta = \min_{i,j} p_{1,j}^i > 0$. We thus have (EC.30) less than or equal to

$$M \sqrt{\frac{8mTN}{\delta R_k} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2}} \quad (\text{EC.31})$$

Therefore, noting that $\psi(\mathbf{p}_k)' \mathbf{d}_k \leq 0$ by the definition of \mathbf{d}_k , from (EC.27) we have

$$E[Z(\mathbf{p}_{k+1}) - Z(\mathbf{p}_k) | \mathcal{F}_k] \leq \epsilon_k M \sqrt{\frac{8mTN}{\delta R_k} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2}} + \frac{\epsilon_k^2 V}{2} \quad (\text{EC.32})$$

and hence

$$\sum_{k=1}^{\infty} E[E[Z(\mathbf{p}_{k+1}) - Z(\mathbf{p}_k) | \mathcal{F}_k]^+] \leq M \sqrt{\frac{8mTN}{\delta} \sum_{k=1}^{\infty} \frac{\epsilon_k}{\sqrt{R_k}} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2}} + \sum_{k=1}^{\infty} \frac{\epsilon_k^2 V}{2}$$

By Assumptions 1 and 2, and Lemma EC.1 (depicted after this proof), we have $Z(\mathbf{p}_k)$ converge to an integrable random variable.

Now take expectation on (EC.27) further to get

$$\begin{aligned} E[Z(\mathbf{p}_{k+1})] &\leq E[Z(\mathbf{p}_k)] + \epsilon_k E[\psi(\mathbf{p}_k)' \mathbf{d}_k] + \epsilon_k E[(\hat{\psi}(\mathbf{p}_k) - \psi(\mathbf{p}_k))' (\mathbf{d}_k - \hat{\mathbf{d}}_k)] \\ &\quad + \frac{\epsilon_k^2}{2} E[\hat{\mathbf{d}}_k' \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k] \end{aligned}$$

and telescope to get

$$\begin{aligned} E[Z(\mathbf{p}_{k+1})] &\leq E[Z(\mathbf{p}_1)] + \sum_{j=1}^k \epsilon_j E[\psi(\mathbf{p}_j)' \mathbf{d}_j] + \sum_{j=1}^k \epsilon_j E[(\hat{\psi}(\mathbf{p}_j) - \psi(\mathbf{p}_j))' (\mathbf{d}_j - \hat{\mathbf{d}}_j)] \\ &\quad + \sum_{j=1}^k \frac{\epsilon_j^2}{2} E[\hat{\mathbf{d}}_j' \nabla^2 Z(\mathbf{p}_j + \theta_j \epsilon_j \hat{\mathbf{d}}_j) \hat{\mathbf{d}}_j] \end{aligned} \quad (\text{EC.33})$$

Now take the limit on both sides of (EC.33). Note that $E[Z(\mathbf{p}_{k+1})] \rightarrow E[Z_\infty]$ for some integrable Z_∞ by dominated convergence theorem. Also $Z(\mathbf{p}_1) < \infty$, and by (EC.28) and (EC.31) respectively, we have

$$\lim_{k \rightarrow \infty} \sum_{j=1}^k \frac{\epsilon_j^2}{2} E[\hat{\mathbf{d}}_j' \nabla^2 Z(\mathbf{p}_j + \theta_j \epsilon_j \hat{\mathbf{d}}_j) \hat{\mathbf{d}}_j] \leq \sum_{j=1}^{\infty} \frac{\epsilon_j^2 V}{2} < \infty$$

and

$$\lim_{k \rightarrow \infty} \sum_{j=1}^k \epsilon_j E[(\hat{\psi}(\mathbf{p}_j) - \psi(\mathbf{p}_j))' (\mathbf{d}_j - \hat{\mathbf{d}}_j)] \leq M \sqrt{\frac{8mTN}{\delta}} \sum_{j=1}^{\infty} \frac{\epsilon_j}{\sqrt{R_j}} \prod_{i=1}^{j-1} (1 - \epsilon_i)^{-1/2} < \infty$$

Therefore, from (EC.33), and since $E[\psi(\mathbf{p}_j)' \mathbf{d}_j] \leq 0$, we must have $\sum_{j=1}^k \epsilon_j E[\psi(\mathbf{p}_j)' \mathbf{d}_j]$ converges a.s., which implies that $\limsup_{k \rightarrow \infty} E[\psi(\mathbf{p}_k)' \mathbf{d}_k] = 0$. So there exists a subsequence k_i such that $\lim_{i \rightarrow \infty} E[\psi(\mathbf{p}_{k_i})' \mathbf{d}_{k_i}] = 0$. This in turn implies that $\psi(\mathbf{p}_{k_i})' \mathbf{d}_{k_i} \xrightarrow{p} 0$. Then, there exists a further subsequence l_i such that $\psi(\mathbf{p}_{l_i})' \mathbf{d}_{l_i} \rightarrow 0$ a.s..

Consider part 1 of the theorem. Let $S^* = \{\mathbf{p} \in \mathcal{P} : g(\mathbf{p}) = 0\}$. Since $g(\cdot)$ is continuous, we have $D(\mathbf{p}_{l_i}, S^*) \rightarrow 0$ a.s.. Since $Z(\cdot)$ is continuous, we have $D(Z(\mathbf{p}_{l_i}), \mathcal{Z}^*) \rightarrow 0$ a.s.. But since we have proven that $Z(\mathbf{p}_k)$ converges a.s., we have $D(Z(\mathbf{p}_k), \mathcal{Z}^*) \rightarrow 0$ a.s.. This gives part 1 of the theorem.

Now consider part 2. By Assumption 3, since \mathbf{p}^* is the only \mathbf{p} such that $g(\mathbf{p}) = 0$ and $g(\cdot)$ is continuous, we must have $\mathbf{p}_{l_i} \rightarrow \mathbf{p}^*$ a.s.. Since $Z(\cdot)$ is continuous, we have $Z(\mathbf{p}_{l_i}) \rightarrow Z(\mathbf{p}^*)$. But since $Z(\mathbf{p}_k)$ converges a.s. as shown above, we must have $Z(\mathbf{p}_k) \rightarrow Z(\mathbf{p}^*)$. Then by Assumption 3 again, since \mathbf{p}^* is the unique optimizer, we have $\mathbf{p}_k \rightarrow \mathbf{p}^*$ a.s.. This concludes part 2 of the theorem.

LEMMA EC.1 (**Adapted from Blum (1954)**). *Consider a sequence of integrable random variable $Y_k, k = 1, 2, \dots$. Let \mathcal{F}_k be the filtration generated by Y_1, \dots, Y_k . Assume*

$$\sum_{k=1}^{\infty} E[E[Y_{k+1} - Y_k | \mathcal{F}_k]^+] < \infty$$

where x^+ denotes the positive part of x , i.e. $x^+ = x$ if $x \geq 0$ and 0 if $x < 0$. Moreover, assume that Y_k is bounded uniformly from above. Then $Y_k \rightarrow Y_{\infty}$ a.s., where Y_{∞} is an integrable random variable.

The lemma follows from Blum (1954), with the additional conclusion that Y_{∞} is integrable, which is a direct consequence of the martingale convergence theorem.

THEOREM EC.1 (Conditions in Theorem 5). *Conditions 1-9 needed in Theorem 5 are:*

1.

$$k_0 \geq 2a \left(\frac{4KMTm}{c^2\tau^2} + \frac{KL\vartheta}{c\tau} \right)$$

2.

$$-\left(1 - \frac{2KL\vartheta}{c\tau} - \frac{2a\varrho K}{c^2\tau^2 k_0}\right)\nu + \frac{2aKL\vartheta\varrho}{c\tau k_0^{1+\gamma}} + \frac{\varrho}{k_0^\gamma} + \frac{2K\nu^2}{c^2\tau^2} \leq 0$$

3.

$$\frac{2KL\vartheta}{c\tau} + \frac{2K\nu}{c^2\tau^2} < 1$$

4.

$$\frac{a}{k_0} \left(1 - \frac{2KL\vartheta}{c\tau} - \frac{2K\nu}{c^2\tau^2}\right) < 1$$

5.

$$k_0 \geq \frac{a\rho}{\rho - 1}$$

6.

$$\beta > \rho a + 2\gamma + 2$$

7.

$$\begin{aligned} & \prod_{j=1}^{k_0-1} (1-\epsilon_j)^{-1} \frac{M^2 TN}{\vartheta^2 \delta b} \frac{1}{(\beta - \rho a - 1)(k_0 - 1)^{\beta-1}} \\ & + \prod_{j=1}^{k_0-1} (1-\epsilon_j)^{-1/2} \frac{M}{\varrho} \sqrt{\frac{8mTN}{\delta b}} \frac{1}{((\beta - \rho a)/2 - \gamma - 1)(k_0 - 1)^{\beta/2-\gamma-1}} < \varepsilon \end{aligned}$$

where $N = \sum_{i=1}^m n^i$ is the total count of all support points.

8. $K > 0$ is a constant such that $|\mathbf{x}' \nabla^2 Z(\mathbf{p}) \mathbf{y}| \leq K \|\mathbf{x}\| \|\mathbf{y}\|$ for any $x, y \in \mathbb{R}^n$ and $\mathbf{p} \in \mathcal{A}$ (which must exist because $Z(\cdot)$ is a polynomial defined over a bounded set).

$$9. \delta = \min_{\substack{i=1, \dots, m \\ j=1, \dots, n^i}} p_{1,j}^i > 0$$

Proof of Theorem 5 We adopt the notation as in the proof of Theorem 4. In addition, for convenience, we write $\psi_k = \psi(\mathbf{p}_k)$, $\hat{\psi}_k = \hat{\psi}(\mathbf{p}_k)$, $\mathbf{q}_k = \mathbf{q}(\mathbf{p}_k)$, $\hat{\mathbf{q}}_k = \hat{\mathbf{q}}(\mathbf{p}_k)$, $g_k = g(\mathbf{p}_k) = -\psi(\mathbf{p}_k)' \mathbf{d}_k$, $\nabla Z_k = \nabla Z(\mathbf{p}_k)$, and $\nabla^2 Z_k = \nabla^2 Z(\mathbf{p}_k)$. Note that $\mathbf{p}_{k+1} = \mathbf{p}_k + \epsilon_k \hat{\mathbf{d}}_k$.

First, by the proof of Theorem 4, given any ν and \tilde{k}_0 , almost surely there must exist a $k_0 \geq \tilde{k}_0$ such that $g_{k_0} \leq \nu$. If the optimal solution is reached and is kept there, then $g_k = 0$ from thereon and the algorithm reaches and remains at optimum at finite time, hence there is nothing to prove. So let us assume that $0 < g_{k_0} \leq \nu$. Moreover, let us assume that ν is chosen small enough so that for any \mathbf{p} with $g(\mathbf{p}) \leq \nu$ and $\mathbf{p} > \mathbf{0}$, we have $\psi(\mathbf{p}) \in \mathcal{N}_{\Delta-\vartheta}(\psi(\mathbf{p}^*))$ (which can be done since $g(\cdot)$ is assumed continuous by Assumption 3 and $\psi(\mathbf{p})$ is continuous for any $\mathbf{p} > \mathbf{0}$ by the construction in Theorem 3).

We consider the event

$$\mathcal{E} = \bigcup_{k=k_0}^{\infty} \mathcal{E}_k \cup \bigcup_{k=k_0}^{\infty} \mathcal{E}'_k$$

where

$$\mathcal{E}_k = \{\|\hat{\psi}_k - \psi_k\| > \vartheta\}$$

and

$$\mathcal{E}'_k = \left\{ |(\hat{\psi}_k - \psi_k)'(\hat{\mathbf{d}}_k - \mathbf{d}_k)| > \frac{\varrho}{k^\gamma} \right\}$$

Note that by the Markov inequality,

$$P(\mathcal{E}_k) \leq \frac{E\|\hat{\psi}_k - \psi_k\|^2}{\vartheta^2} \leq \frac{M^2 T}{\vartheta^2 R_k} \sum_{i,j} \frac{1 - p_{k,j}^i}{p_{k,j}^i} \leq \frac{M^2 T N}{\vartheta^2 R_k \delta} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1}$$

where the second inequality follows from Lemma 1 and the last inequality follows as in the derivation in (EC.30) and (EC.31). On the other hand, we have

$$P(\mathcal{E}'_k) \leq \frac{k^\gamma E|(\hat{\psi}_k - \psi_k)'(\hat{\mathbf{d}}_k - \mathbf{d}_k)|}{\varrho} \leq \frac{k^\gamma M}{\varrho} \sqrt{\frac{8mTN}{\delta R_k}} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2} \quad (\text{EC.34})$$

by following the derivation in (EC.30) and (EC.31). Therefore,

$$\begin{aligned} P(\mathcal{E}) &\leq \sum_{k=k_0}^{\infty} P(\mathcal{E}_k) + \sum_{k=k_0}^{\infty} P(\mathcal{E}'_k) \\ &\leq \frac{M^2 TN}{\vartheta^2 \delta} \sum_{k=k_0}^{\infty} \frac{1}{R_k} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1} + \frac{M}{\varrho} \sqrt{\frac{8mTN}{\delta}} \sum_{k=k_0}^{\infty} \frac{k^\gamma}{\sqrt{R_k}} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1/2} \\ &= \prod_{j=1}^{k_0-1} (1 - \epsilon_j)^{-1} \frac{M^2 TN}{\vartheta^2 \delta} \sum_{k=k_0}^{\infty} \frac{1}{R_k} \prod_{j=k_0}^{k-1} (1 - \epsilon_j)^{-1} + \prod_{j=1}^{k_0-1} (1 - \epsilon_j)^{-1/2} \frac{M}{\varrho} \sqrt{\frac{8mTN}{\delta}} \sum_{k=k_0}^{\infty} \frac{k^\gamma}{\sqrt{R_k}} \prod_{j=k_0}^{k-1} (1 - \epsilon_j)^{-1/2} \end{aligned} \quad (\text{EC.35})$$

Now recall that $\epsilon_k = a/k$. Using the fact that $1 - x \geq e^{-\rho x}$ for any $0 \leq x \leq (\rho - 1)/\rho$ and $\rho > 1$, we have, for any

$$\frac{a}{k} \leq \frac{\rho - 1}{\rho}$$

or equivalently

$$k \geq \frac{a\rho}{\rho - 1}$$

we have

$$1 - \epsilon_k = 1 - \frac{a}{k} \geq e^{-\rho a/k}$$

Hence choosing k_0 satisfying Condition 5, we get

$$\prod_{j=k_0}^{k-1} (1 - \epsilon_j)^{-1} \leq e^{\rho a \sum_{j=k_0}^{k-1} 1/j} \leq \left(\frac{k-1}{k_0-1} \right)^{\rho a} \quad (\text{EC.36})$$

Therefore, picking $R_k = bk^\beta$ and using (EC.36), we have (EC.35) bounded from above by

$$\begin{aligned} & \prod_{j=1}^{k_0-1} (1-\epsilon_j)^{-1} \frac{M^2 TN}{\vartheta^2 \delta b} \sum_{k=k_0}^{\infty} \frac{1}{(k_0-1)^{\rho a} k^{\beta-\rho a}} + \prod_{j=1}^{k_0-1} (1-\epsilon_j)^{-1/2} \frac{M}{\varrho} \sqrt{\frac{8mTN}{\delta b}} \sum_{k=k_0}^{\infty} \frac{1}{(k_0-1)^{\rho a/2} k^{(\beta-\rho a)/2-\gamma}} \\ & \leq \prod_{j=1}^{k_0-1} (1-\epsilon_j)^{-1} \frac{M^2 TN}{\vartheta^2 \delta b} \frac{1}{(\beta-\rho a-1)(k_0-1)^{\beta-1}} \\ & \quad + \prod_{j=1}^{k_0-1} (1-\epsilon_j)^{-1/2} \frac{M}{\varrho} \sqrt{\frac{8mTN}{\delta b}} \frac{1}{((\beta-\rho a)/2-\gamma-1)(k_0-1)^{\beta/2-\gamma-1}} \end{aligned} \quad (\text{EC.37})$$

if Condition 6 holds. Then Condition 7 guarantees that $P(\mathcal{E}) < \varepsilon$.

The rest of the proof will show that under the event \mathcal{E}^c , we must have the bound (26), hence concluding the theorem. To this end, we first set up a recursive representation of g_k . Consider

$$\begin{aligned} g_{k+1} &= -\psi'_{k+1} \mathbf{d}_{k+1} = -\psi'_{k+1} (\mathbf{q}_{k+1} - \mathbf{p}_{k+1}) \\ &= -\psi'_k (\mathbf{q}_{k+1} - \mathbf{p}_{k+1}) + (\psi_k - \psi_{k+1})' (\mathbf{q}_{k+1} - \mathbf{p}_{k+1}) \\ &= -\psi'_k (\mathbf{q}_{k+1} - \mathbf{p}_k) + \psi'_k (\mathbf{p}_{k+1} - \mathbf{p}_k) + (\psi_k - \psi_{k+1})' (\mathbf{q}_{k+1} - \mathbf{p}_{k+1}) \\ &\leq g_k + \epsilon_k \psi'_k \hat{\mathbf{d}}_k + (\psi_k - \psi_{k+1})' \mathbf{d}_{k+1} \quad \text{by the definition of } g_k, \hat{\mathbf{d}}_k \text{ and } \mathbf{d}_{k+1} \\ &\leq g_k - \epsilon_k g_k + \epsilon_k (\hat{\psi}_k - \psi_k)' (\mathbf{d}_k - \hat{\mathbf{d}}_k) + (\psi_k - \psi_{k+1})' \mathbf{d}_{k+1} \quad \text{by (EC.26)} \\ &= (1 - \epsilon_k) g_k + (\nabla Z_k - \nabla Z_{k+1})' \mathbf{d}_{k+1} + \epsilon_k (\hat{\psi}_k - \psi_k)' (\mathbf{d}_k - \hat{\mathbf{d}}_k) \end{aligned} \quad (\text{EC.38})$$

Now since $\nabla Z(\cdot)$ is continuously differentiable, we have $\nabla Z_{k+1} = \nabla Z_k + \epsilon_k \nabla^2 Z(\mathbf{p}_k + \tilde{\theta}_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k$ for some $\tilde{\theta}_k$ between 0 and 1. Therefore (EC.38) is equal to

$$\begin{aligned} & (1 - \epsilon_k) g_k - \epsilon_k \hat{\mathbf{d}}'_k \nabla^2 Z(\mathbf{p}_k + \tilde{\theta}_k \hat{\mathbf{d}}_k) \mathbf{d}_{k+1} + \epsilon_k (\hat{\psi}_k - \psi_k)' (\mathbf{d}_k - \hat{\mathbf{d}}_k) \\ & \leq (1 - \epsilon_k) g_k + \epsilon_k K \|\hat{\mathbf{d}}_k\| \|\mathbf{d}_{k+1}\| + \epsilon_k (\hat{\psi}_k - \psi_k)' (\mathbf{d}_k - \hat{\mathbf{d}}_k) \quad \text{by Condition 8} \\ & \leq (1 - \epsilon_k) g_k + \epsilon_k K \|\mathbf{d}_k\| \|\mathbf{d}_{k+1}\| + \epsilon_k K \|\hat{\mathbf{d}}_k - \mathbf{d}_k\| \|\mathbf{d}_{k+1}\| + \epsilon_k (\hat{\psi}_k - \psi_k)' (\mathbf{d}_k - \hat{\mathbf{d}}_k) \end{aligned}$$

by the triangle inequality

$$\leq (1 - \epsilon_k) g_k + \epsilon_k K \frac{g_k g_{k+1}}{c^2 \|\psi_k\| \|\psi_{k+1}\|} + \epsilon_k K L \|\hat{\psi}_k - \psi_k\| \frac{g_{k+1}}{c \|\psi_{k+1}\|} + \epsilon_k (\hat{\psi}_k - \psi_k)' (\mathbf{d}_k - \hat{\mathbf{d}}_k)$$

by using Assumption 4 with the fact that $g_k \leq \nu$ and hence $\psi_k, \hat{\psi}_k \in \mathcal{N}_\Delta(\psi(\mathbf{p}^*))$, and also

Assumption 5. The fact $g_k \leq \nu$ will be proved later by induction.

$$\leq (1 - \epsilon_k) g_k + \epsilon_k \frac{K}{c^2 \tau^2} g_k g_{k+1} + \epsilon_k \frac{KL}{c\tau} \|\hat{\psi}_k - \psi_k\| g_{k+1} + \epsilon_k (\hat{\psi}_k - \psi_k)' (\mathbf{d}_k - \hat{\mathbf{d}}_k) \quad (\text{EC.39})$$

by Assumption 6

Now under the event \mathcal{E}^c , and noting that $\epsilon = a/k$, (EC.39) implies that

$$g_{k+1} \leq \left(1 - \frac{a}{k}\right) g_k + \frac{aK}{c^2\tau^2 k} g_k g_{k+1} + \frac{aKL\vartheta}{c\tau k} g_{k+1} + \frac{a\varrho}{k^{1+\gamma}}$$

or

$$\left(1 - \frac{aK}{c^2\tau^2 k} g_k - \frac{aKL\vartheta}{c\tau k}\right) g_{k+1} \leq \left(1 - \frac{a}{k}\right) g_k + \frac{a\varrho}{k^{1+\gamma}}$$

We claim that $|g_k| = |\psi'_k \mathbf{d}_k| \leq 4MTm$, which can be seen by writing

$$\begin{aligned} \psi_j^i(\mathbf{p}) &= E_{\mathbf{p}}[h(\mathbf{X})s_j^i(\mathbf{X}^i)] = \sum_{t=1}^{T^i} E_{\mathbf{p}}\left[h(\mathbf{X}) \frac{I(X_t^i = y_j^i)}{p_j^i}\right] - T^i E_{\mathbf{p}}[h(\mathbf{X})] \\ &= \sum_{t=1}^{T^i} E_{\mathbf{p}}[h(\mathbf{X})|X_t = y_j^i] - T^i E_{\mathbf{p}}[h(\mathbf{X})] \end{aligned} \quad (\text{EC.40})$$

so that $|\psi_j^i(\mathbf{p})| \leq 2MT^i$ for any \mathbf{p} and i . Using this and the fact that $1/(1-x) \leq 1 + 2x$ for any $0 \leq x \leq 1/2$, we have, for

$$\frac{4aKMTm}{c^2\tau^2 k} + \frac{aKL\vartheta}{c\tau k} \leq \frac{1}{2} \quad (\text{EC.41})$$

we must have

$$g_{k+1} \leq \left(1 + \frac{2aK}{c^2\tau^2 k} g_k + \frac{2aKL\vartheta}{c\tau k}\right) \left(\left(1 - \frac{a}{k}\right) g_k + \frac{a\varrho}{k^{1+\gamma}}\right) \quad (\text{EC.42})$$

Note that (EC.41) holds if

$$k \geq 2a \left(\frac{4KMTm}{c^2\tau^2} + \frac{KL\vartheta}{c\tau} \right)$$

which is Condition 1 in the theorem. Now (EC.42) can be written as

$$\begin{aligned} g_{k+1} &\leq \left(1 - \frac{a}{k} + \frac{2aKL\vartheta}{c\tau k} + \frac{2a^2K\varrho}{c^2\tau^2 k^{2+\gamma}}\right) g_k + \frac{a\varrho}{k^{1+\gamma}} + \frac{2a^2KL\vartheta\varrho}{c\tau k^{2+\gamma}} - \frac{2a^2KL\vartheta}{c\tau k^2} g_k + \frac{2aK}{c^2\tau^2 k} \left(1 - \frac{a}{k}\right) g_k^2 \\ &\leq \left(1 - \frac{a}{k} + \frac{2aKL\vartheta}{c\tau k} + \frac{2a^2K\varrho}{c^2\tau^2 k^{2+\gamma}}\right) g_k + \frac{a\varrho}{k^{1+\gamma}} + \frac{2a^2KL\vartheta\varrho}{c\tau k^{2+\gamma}} + \frac{2aK}{c^2\tau^2 k} \left(1 - \frac{a}{k}\right) g_k^2 \end{aligned} \quad (\text{EC.43})$$

We argue that under Condition 2, we must have $g_k \leq \nu$ for all $k \geq k_0$. This can be seen by induction using (EC.43). By our setting at the beginning of this proof we have $g_{k_0} \leq \nu$. Suppose $g_k \leq \nu$ for some k . We then have

$$\begin{aligned} g_{k+1} &\leq \left(1 - \frac{a}{k} + \frac{2aKL\vartheta}{c\tau k} + \frac{2a^2K\varrho}{c^2\tau^2 k^{2+\gamma}}\right) \nu + \frac{a\varrho}{k^{1+\gamma}} + \frac{2a^2KL\vartheta\varrho}{c\tau k^{2+\gamma}} + \frac{2aK}{c^2\tau^2 k} \left(1 - \frac{a}{k}\right) \nu^2 \\ &\leq \nu + \frac{a}{k} \left(\left(-1 + \frac{2KL\vartheta}{c\tau} + \frac{2aK\varrho}{c^2\tau^2 k^{1+\gamma}}\right) \nu + \frac{\varrho}{k_0^\gamma} + \frac{2aKL\vartheta\varrho}{c\tau k_0^{1+\gamma}} + \frac{2K\nu^2}{c^2\tau^2} \right) \\ &\leq \nu \end{aligned} \quad (\text{EC.44})$$

by Condition 2. This concludes our claim.

Given that $g_k \leq \nu$ for all $k \geq k_0$, (EC.42) implies that

$$\begin{aligned} g_{k+1} &\leq \left(1 - \frac{a}{k} \left(1 - \frac{2KL\vartheta}{c\tau}\right) - \frac{2a^2KL\vartheta}{c\tau k^2} + \frac{2aK\nu}{c^2\tau^2 k} \left(1 - \frac{a}{k}\right)\right) g_k + \frac{a\varrho}{k^{1+\gamma}} + \frac{a^2\varrho}{k^{2+\gamma}} \left(\frac{2K\nu}{c^2\tau^2} + \frac{2KL\vartheta}{c\tau}\right) \\ &\leq \left(1 - \frac{a}{k} \left(1 - \frac{2KL\vartheta}{c\tau} - \frac{2K\nu}{c^2\tau^2}\right)\right) g_k + \frac{a\varrho}{k^{1+\gamma}} + \frac{a^2\varrho}{k^{2+\gamma}} \left(\frac{2K\nu}{c^2\tau^2} + \frac{2KL\vartheta}{c\tau}\right) \\ &\leq \left(1 - \frac{C}{k}\right) g_k + \frac{G}{k^{1+\gamma}} \end{aligned} \quad (\text{EC.45})$$

where

$$C = a \left(1 - \frac{2KL\vartheta}{c\tau} - \frac{2K\nu}{c^2\tau^2}\right)$$

and

$$G = a\varrho + \frac{a^2\varrho}{k_0} \left(\frac{2K\nu}{c^2\tau^2} + \frac{2KL\vartheta}{c\tau}\right)$$

Now note that Conditions 3 and 4 imply $C > 0$ and $1 - C/k > 0$ respectively. By recursing the relation (EC.45), we get

$$\begin{aligned} g_{k+1} &\leq \prod_{j=k_0}^k \left(1 - \frac{C}{j}\right) g_{k_0} + \sum_{j=k_0}^k \prod_{i=j+1}^k \left(1 - \frac{C}{i}\right) \frac{G}{j^{1+\gamma}} \\ &\leq e^{-C \sum_{j=k_0}^k 1/j} g_{k_0} + \sum_{j=k_0}^k e^{-C \sum_{i=j+1}^k 1/i} \frac{G}{j^{1+\gamma}} \\ &\leq \left(\frac{k_0}{k+1}\right)^C g_{k_0} + \sum_{j=k_0}^k \left(\frac{j+1}{k+1}\right)^C \frac{G}{j^{1+\gamma}} \\ &\leq \left(\frac{k_0}{k+1}\right)^C g_{k_0} + \left(1 + \frac{1}{k_0}\right)^C G \times \begin{cases} \frac{1}{(C-\gamma)(k+1)^\gamma} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^{\gamma-C}(k+1)^C} & \text{if } \gamma > C \\ \frac{\log(k/(k_0-1))}{(k+1)^C} & \text{if } \gamma = C \end{cases} \end{aligned}$$

which gives (26). This concludes the proof.

Proof of Corollary 1 We use the notations in the proof of Theorem 5. Our analysis starts from (EC.25), namely

$$Z_{k+1} = Z_k + \epsilon_k \psi'_k \hat{\mathbf{d}}_k + \frac{\epsilon_k^2}{2} \hat{\mathbf{d}}'_k \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k$$

for some $\tilde{\theta}_k$ between 0 and 1. Using the fact that $\psi'_k \hat{\mathbf{d}}_k \geq \psi'_k \mathbf{d}_k$ by the definition of \mathbf{d}_k , we have

$$\begin{aligned} Z_{k+1} &\geq Z_k + \epsilon_k \psi'_k \mathbf{d}_k + \frac{\epsilon_k^2}{2} \hat{\mathbf{d}}'_k \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k \\ &= Z_k - \epsilon_k g_k + \frac{\epsilon_k^2}{2} \hat{\mathbf{d}}'_k \nabla^2 Z(\mathbf{p}_k + \theta_k \epsilon_k \hat{\mathbf{d}}_k) \hat{\mathbf{d}}_k \end{aligned}$$

Now, using (26), Condition 8 in Theorem 5 and $\|\hat{\mathbf{d}}_k\|^2 \leq 2$, we have

$$\begin{aligned} Z_{k+1} &\geq Z_k - \epsilon_k \left(\frac{A}{k^C} + B \times \begin{cases} \frac{1}{(C-\gamma)k^\gamma} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^{\gamma-C} k^C} & \text{if } \gamma > C \\ \frac{\log((k-1)/(k_0-1))}{k^C} & \text{if } \gamma = C \end{cases} \right) - \epsilon_k^2 K \\ &= Z_k - \frac{aA}{k^{1+C}} - aB \times \left(\begin{cases} \frac{1}{(C-\gamma)k^{1+\gamma}} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^{\gamma-C} k^{1+C}} & \text{if } \gamma > C \\ \frac{\log((k-1)/(k_0-1))}{k^{1+C}} & \text{if } \gamma = C \end{cases} \right) - \frac{a^2 K}{k^2} \end{aligned} \quad (\text{EC.46})$$

Now iterating (EC.46) from k to l , we have

$$Z_l \geq Z_k - \sum_{j=k}^{l-1} \frac{aA}{j^{1+C}} - aB \times \left(\begin{cases} \frac{1}{(C-\gamma)} \sum_{j=k}^{l-1} \frac{1}{j^{1+\gamma}} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^{\gamma-C}} \sum_{j=k}^{l-1} \frac{1}{j^{1+C}} & \text{if } \gamma > C \\ \sum_{j=k}^{l-1} \frac{\log((j-1)/(k_0-1))}{j^{1+C}} & \text{if } \gamma = C \end{cases} \right) - a^2 K \sum_{j=k}^{l-1} \frac{1}{j^2}$$

and letting $l \rightarrow \infty$, we get

$$Z^* \geq Z_k - \sum_{j=k}^{\infty} \frac{aA}{j^{1+C}} - aB \times \left(\begin{cases} \frac{1}{(C-\gamma)} \sum_{j=k}^{\infty} \frac{1}{j^{1+\gamma}} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^{\gamma-C}} \sum_{j=k}^{\infty} \frac{1}{j^{1+C}} & \text{if } \gamma > C \\ \sum_{j=k}^{\infty} \frac{\log((j-1)/(k_0-1))}{j^{1+C}} & \text{if } \gamma = C \end{cases} \right) - a^2 K \sum_{j=k}^{\infty} \frac{1}{j^2} \quad (\text{EC.47})$$

where the convergence to Z^* is guaranteed by Theorem 4. Note that (EC.47) implies that

$$\begin{aligned} Z^* &\geq Z_k - \frac{aA}{C(k-1)^C} - aB \times \left(\begin{cases} \frac{1}{(C-\gamma)\gamma(k-1)^\gamma} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^{\gamma-C} C(k-1)^C} & \text{if } \gamma > C \\ \frac{\log((k-1)/(k_0-1))}{C(k-1)^C} & \text{if } \gamma = C \end{cases} \right) - \frac{a^2 K}{k-1} \\ &\geq Z_k - \frac{D}{k-1} - \frac{E}{(k-1)^C} - F \times \left(\begin{cases} \frac{1}{(C-\gamma)\gamma(k-1)^\gamma} & \text{if } 0 < \gamma < C \\ \frac{1}{(\gamma-C)(k_0-1)^{\gamma-C} C(k-1)^C} & \text{if } \gamma > C \\ \frac{\log((k-1)/(k_0-1))}{C(k-1)^C} & \text{if } \gamma = C \end{cases} \right) \end{aligned}$$

where $D = a^2 K$, $E = aA/C$ and $F = aB$. This gives (28).

Proof of Lemma 2 Consider first a fixed a . When $a(1 - \omega) > 1$, (29) reduces to $\frac{\beta - \rho a - \zeta - 2}{2(\beta + 1)} \wedge \frac{1}{\beta + 1}$.

Since $\frac{\beta - \rho a - \zeta - 2}{2(\beta + 1)}$ is increasing in β and $\frac{1}{\beta + 1}$ is decreasing in β , the maximizer of $\frac{\beta - \rho a - \zeta - 2}{2(\beta + 1)} \wedge \frac{1}{\beta + 1}$ occurs at the intersection of $\frac{\beta - \rho a - \zeta - 2}{2(\beta + 1)}$ and $\frac{1}{\beta + 1}$, which is $\beta = \rho a + \zeta + 4$. The associated value of (29) is $\frac{1}{\rho a + \zeta + 5}$.

When $a(1 - \omega) \leq 1$, (29) reduces to $\frac{a(1 - \omega)}{\beta + 1} \wedge \frac{\beta - \rho a - \zeta - 2}{2(\beta + 1)}$. By a similar argument, the maximizer is $\beta = a(2 - 2\omega + \rho) + \zeta + 2$, with the value of (29) equal to $\frac{a(1 - \omega)}{a(2 - 2\omega + \rho) + \zeta + 3}$.

Thus, overall, given a , the optimal choice of β is $\beta = \rho a + \zeta + 2 + 2((a(1 - \omega)) \wedge 1)$, with the value of (29) given by $\frac{(a(1 - \omega)) \wedge 1}{\rho a + \zeta + 3 + 2((a(1 - \omega)) \wedge 1)}$. When $a(1 - \omega) > 1$, the value of (29) is $\frac{1}{\rho a + \zeta + 5}$ which is decreasing in a , whereas when $a(1 - \omega) \leq 1$, the value of (29) is $\frac{a(1 - \omega)}{a(2 - 2\omega + \rho) + \zeta + 3}$ which is increasing in a . Thus the maximum occurs when $a(1 - \omega) = 1$, or $a = \frac{1}{1 - \omega}$. The associated value of (29) is $\frac{1}{\rho/(1 - \omega) + \zeta + 5}$.

REMARK EC.1. Suppose that Assumption 4 is replaced by letting

$$\|\mathbf{v}(\boldsymbol{\xi}_1) - \mathbf{v}(\boldsymbol{\xi}_2)\| \leq L\|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|$$

hold for any $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2 \in \mathbb{R}^N$. Then, in the proof of Theorem 5, the inequality (EC.34) can be replaced by

$$\begin{aligned} P(\mathcal{E}'_k) &\leq \frac{k^\gamma E|(\hat{\boldsymbol{\psi}}_k - \boldsymbol{\psi}_k)'(\hat{\mathbf{d}}_k - \mathbf{d}_k)|}{\varrho} \\ &\leq \frac{k^\gamma}{\varrho} \sqrt{E[\|\hat{\boldsymbol{\psi}}_k - \boldsymbol{\psi}_k\|^2]E[\|\mathbf{d}_k - \hat{\mathbf{d}}_k\|^2]} \quad \text{by the Cauchy-Schwarz inequality} \\ &\leq \frac{k^\gamma L}{\varrho} E[\|\hat{\boldsymbol{\psi}}_k - \boldsymbol{\psi}_k\|^2] \quad \text{by the relaxed Assumption 4} \\ &\leq \frac{LM^2TNk^\gamma}{R_k\varrho\delta} \prod_{j=1}^{k-1} (1 - \epsilon_j)^{-1} \quad \text{by following the derivation in (EC.30) and (EC.31)} \end{aligned}$$

Consequently, equation (EC.37) becomes

$$\prod_{j=1}^{k_0-1} (1 - \epsilon_j)^{-1} \frac{M^2TN}{\delta b} \left(\frac{1}{\vartheta^2(\beta - \rho a - 1)(k_0 - 1)^{\beta-1}} + \frac{L}{\varrho(\beta - \gamma - \rho a - 1)(k_0 - 1)^{\beta-\gamma-1}} \right)$$

if Condition 6 is replaced by

$$\beta > \gamma + \rho a + 1$$

Correspondingly, Condition 7 needs to be replaced by

$$\prod_{j=1}^{k_0-1} (1-\epsilon_j)^{-1} \frac{M^2 TN}{\delta b} \left(\frac{1}{\vartheta^2(\beta-\rho a-1)(k_0-1)^{\beta-1}} + \frac{L}{\varrho(\beta-\gamma-\rho a-1)(k_0-1)^{\beta-\gamma-1}} \right) < \varepsilon$$

The results in Theorem 5 and Corollary 1 then retain. Under these modified Conditions 6 and 7, discussion point 3(b) in Section 6.2 then gives $\beta = \gamma + \rho a + 1 + \zeta$ for some $\zeta > 0$ and $\gamma = \beta - \rho a - \zeta - 1$. In discussion point 4, the convergence rate in terms of replications becomes $1/W^{((a(1-\omega)) \wedge (\beta-\rho a-\zeta-1) \wedge 1) / (\beta+1)}$. By maximizing

$$\frac{(a(1-\omega)) \wedge (\beta-\rho a-\zeta-1) \wedge 1}{\beta+1} \quad (\text{EC.48})$$

like in (29) by Lemma 2 (see Lemma EC.2 right after this remark), we get

$$a = \frac{1}{1-\omega}, \quad \beta = \frac{\rho}{1-\omega} + \zeta + 2$$

and the optimal value is

$$\frac{1}{\rho/(1-\omega) + \zeta + 3}$$

So, following the argument there, we choose ϑ and ν , and hence ω , to be small, and we choose ρ to be close to 1. This gives rise to the approximate choice that $a \approx 1 + \omega$ and $\beta \approx 3 + \zeta + \omega$. The convergence rate is then $O(W^{-1/(4+\zeta+\omega)})$, leading to our claim in Section 6.2 that the complexity can improve to $O(1/\epsilon^{4+\zeta+\omega})$ if Assumption 4 is relaxed.

LEMMA EC.2. *The maximizer of (EC.48) is given by*

$$a = \frac{1}{1-\omega}, \quad \beta = \frac{\rho}{1-\omega} + \zeta + 2$$

and the optimal value is

$$\frac{1}{\rho/(1-\omega) + \zeta + 3}$$

Proof of Lemma EC.2 Consider first a fixed a . When $a(1-\omega) > 1$, (EC.48) reduces to $\frac{\beta-\rho a-\zeta-1}{\beta+1} \wedge \frac{1}{\beta+1}$. Since $\frac{\beta-\rho a-\zeta-1}{\beta+1}$ is increasing in β and $\frac{1}{\beta+1}$ is decreasing in β , the maximizer of

$\frac{\beta-\rho a-\zeta-1}{\beta+1} \wedge \frac{1}{\beta+1}$ occurs at the intersection of $\frac{\beta-\rho a-\zeta-1}{\beta+1}$ and $\frac{1}{\beta+1}$, which is $\beta = \rho a + \zeta + 2$. The associated value of (EC.48) is $\frac{1}{\rho a + \zeta + 3}$.

When $a(1-\omega) \leq 1$, (EC.48) reduces to $\frac{a(1-\omega)}{\beta+1} \wedge \frac{\beta-\rho a-\zeta-1}{\beta+1}$. By a similar argument, the maximizer is $\beta = a(1-\omega+\rho) + \zeta + 1$, with the value of (EC.48) equal to $\frac{a(1-\omega)}{a(1-\omega+\rho)+\zeta+2}$.

Thus, overall, given a , the optimal choice of β is $\beta = \rho a + \zeta + 1 + (a(1-\omega)) \wedge 1$, with the value of (EC.48) given by $\frac{(a(1-\omega)) \wedge 1}{\rho a + \zeta + 2 + (a(1-\omega)) \wedge 1}$. When $a(1-\omega) > 1$, the value of (EC.48) is $\frac{1}{\rho a + \zeta + 3}$ which is decreasing in a , whereas when $a(1-\omega) \leq 1$, the value of (EC.48) is $\frac{a(1-\omega)}{a(1-\omega+\rho)+\zeta+2}$ which is increasing in a . Thus the maximum occurs when $a(1-\omega) = 1$, or $a = \frac{1}{1-\omega}$. The associated value of (EC.48) is $\frac{1}{\rho/(1-\omega)+\zeta+3}$.

EC.2. Additional Details of the Numerical Results

EC.2.1. Multi-start Initialization

The results in Section 7.1 are implemented with an initialization that assigns equal probabilities to the support points. To test the procedure under different initializations, we repeat ten runs of the FWSA algorithm where the initial probability masses for the support points (held constant for all runs) are sampled uniformly independently with appropriate normalization. Figure EC.1 provides a box-plot of the identified optima. The sample size for moment constraint generation is $N_s = 50$ and the discretization support size is $n = 30$. The returned optimal solutions for each of the minimization and maximization formulations all agree up to the first two digits (the box plot shows the small spread of the max values, while the min values are very clustered and they appear to all overlap at the same point). This indicates that the formulations have a unique global optimal solution or similar local optimal solutions. Note that the bounds generated from this setting are quite loose with a small N_s .

EC.2.2. Details of the Benchmark Steady-State Formulation in Section 7.2

We consider the depicted $Z(\mathbf{p})$ in Section 7.2. As T grows, the average waiting time converges to the corresponding steady-state value, which, when the traffic intensity $\rho_{\mathbf{p}} = E_{\mathbf{p}}[X_t]$ is less than 1,



Figure EC.1 Returned optimal solutions from 10 runs on $n = 30$, $M = 50$, exponential for discretization

is given in closed form by the Pollaczek-Khinchine formula (Khintchine (1932)) as:

$$Z_\infty(\mathbf{p}) = \frac{\rho_{\mathbf{p}} E_{\mathbf{p}}[X_1] + \text{Var}_{\mathbf{p}}(X_1)}{2(1 - \rho_{\mathbf{p}})}.$$

So, when T is large, an approximation Z_∞^* to the worst-case performance estimate can be obtained by replacing $Z(\mathbf{p})$ with $Z_\infty(\mathbf{p})$. (In experiments, a choice of $T = 500$ seems to show close agreement.) With $E_{\mathbf{p}}[X_1] = \sum p_j y_j$ and $E_{\mathbf{p}}[X_1^2] = \sum p_j y_j^2$, the steady-state approximation to (32) is given by (SS) below, which is equivalent to (SS') via variable substitutions (see p.191 in Boyd and Vandenberghe (2009)):

$$\begin{aligned}
 \min_{\mathbf{p}} \quad & \frac{\sum_j p_j y_j^2}{2(1 - \sum_j p_j y_j)} && \text{(SS)} \\
 \text{s.t.} \quad & \sum_j p_j \log\left(\frac{p_j}{p_{b,j}}\right) \leq \eta t \\
 & \sum_j p_j = 1 \\
 & 0 \leq p_j \leq 1, \quad \forall j = 1, \dots, n
 \end{aligned}
 \implies
 \begin{aligned}
 \min_{\mathbf{p}} \quad & \sum_j w_j y_j^2 && \text{(SS')} \\
 \text{s.t.} \quad & \sum_j w_j \log\left(\frac{w_j}{t p_{b,j}}\right) \leq \eta t \\
 & 2t - 2 \sum_j w_j y_j = 1 \\
 & \sum_j w_j = t \\
 & 0 \leq w_j \leq t \quad \forall j = 1, \dots, n
 \end{aligned}$$

EC.2.3. Shape of the Obtained Optimal Distributions in Section 7.2

Continuing with the example in Section 7.2, Figure EC.2 shows the form of the optimal distributions \mathbf{p}^* identified by the FWSA algorithm for the minimization (Figure EC.2a) and maximization (Figure EC.2b) problems under (32). The optimal distributions follow a similar bimodal structure as the baseline distribution \mathbf{p}_b . The maximization version assigns probability masses in an unequal manner to the two modes in order to drive up both the mean and the variance of \mathbf{p} , as (SS)

(in Appendix EC.2.2) leads us to expect, whereas the minimization version on the other hand makes the mass allocation more equal in order to minimize the mean and the variance of \mathbf{p} while maintaining the maximum allowed KL divergence.

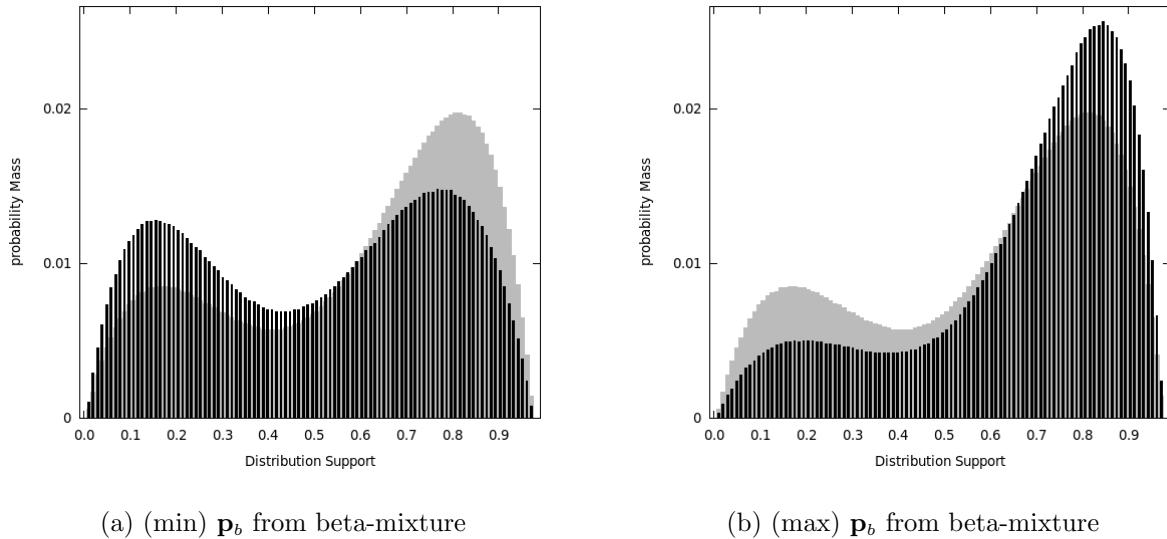


Figure EC.2 Optimal solutions \mathbf{p}^* identified by the FWSA algorithm with $n = 100$ and $\eta = 0.05$, setting $a = 1.5, \beta = 2.75$. The gray bars represent the baseline p.m.f. \mathbf{p}_b .