Infinite Swapping using IID Samples

PAUL DUPUIS, GUO-JHEN WU, and MICHAEL SNARSKI, Division of Applied Mathematics, Brown University, Providence, Rhode Island

We propose a new method for estimating rare event probabilities when independent samples are available. It is assumed that the underlying probability measures satisfy a large deviation principle with a scaling parameter ε that we call temperature. We show how by combining samples at different temperatures, one can construct an estimator with greatly reduced variance. Although as presented here the method is not as broadly applicable as other rare event simulation methods, such as splitting or importance sampling, it does not require any problem-dependent constructions.

CCS Concepts: • **Mathematics of computing** → **Probabilistic algorithms**; • **Computing methodologies** → **Rare-event simulation**;

Additional Key Words and Phrases: Large deviation principle, Monte Carlo method, parallel tempering, rare event simulation

ACM Reference format:

Paul Dupuis, Guo-Jhen Wu, and Michael Snarski. 2019. Infinite Swapping using IID Samples. *ACM Trans. Model. Comput. Simul.* 29, 3, Article 13 (June 2019), 26 pages.
<https://doi.org/10.1145/3317605>

1 INTRODUCTION

The purpose of this article is to introduce a new method for estimating probabilities of rare events and expected values that are largely determined by rare events. In the form developed here, the method makes significant assumptions on the probability measures involved, and for this reason is less broadly applicable than other methods such as importance sampling and splitting. However, when applicable, it is remarkably easy to use, and gives large variance reduction while requiring very little in the way of preliminary constructions, such as importance functions.

The method assumes one can sample directly from the target distribution and certain related distributions. In particular, the generation of samples is not based on the simulation of a Markov chain. However, it is directly motivated by a scheme known as *infinite swapping* (INS) that is applied to systems sampled using Markov-chain Monte Carlo (MCMC) and for which one does not sample directly from the target distribution. Although the focus of this article will be to introduce and optimize the scheme in a large deviation limit when direct sampling is possible, the lessons

Research was supported in part by DOE DE-SC-0010539 and DARPA EQUiPS W911NF-15-2-0122, NSF DMS-1317199 (PD), and NSERC PGSD3-471576-2015 (MS).

Authors' address: P. Dupuis, G.-J. Wu, and M. Snarski, Division of Applied Mathematics, Brown University, 182 George Street, Providence, RI, 02912; emails: {paul_dupuis, guo-jhen_wu}@brown.edu, michael_snarski@alumni.brown.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

1049-3301/2019/06-ART13 \$15.00

<https://doi.org/10.1145/3317605>

learned here can be applied to the more generally applicable INS scheme, and this will be a topic for future work.

An outline of the article is as follows: In the next section, we state assumptions, specify the estimation problem of interest, and introduce the estimator. Section 3 analyzes the performance of the scheme in the large deviation limit by characterizing the decay rate of the second moment of the estimator as a function of certain design parameters, and then shows how these parameters can be selected to optimize the decay rate. Section 4 presents generalizations of the framework, and computational examples are given in Section 5. The method has certain features in common with schemes often used for these problems (importance sampling and multi-level splitting), and a discussion on these points is given in Section 6. We end with a section of concluding remarks. An appendix which gives proofs of various supporting results can be found in the online supplementary material.

2 PROBLEM FORMULATION AND DEFINITION OF THE ESTIMATOR

2.1 Problems of Interest

Consider a collection of probability measures $\{\mu^\varepsilon\}_{\varepsilon \in (0, \infty)}$ on a Polish space $\mathcal{X}$ (i.e., a complete separable metric space). A rate function is a mapping $I : \mathcal{X} \rightarrow [0, \infty]$ such that $\{x : I(x) \leq M\}$ is compact for each $M \in [0, \infty)$. We assume $\{\mu^\varepsilon\}$ satisfies the following properties:

- (1) $\{\mu^\varepsilon\}$ satisfies a large deviation principle (LDP) [3] with rate function I that is continuous on its domain of finiteness and scaling ε ,
- (2) The measures $\mu^{\varepsilon/\alpha_1}$ and $\mu^{\varepsilon/\alpha_2}$ are mutually absolutely continuous for $\alpha_1, \alpha_2 \in (0, 1]$, and we can explicitly compute quantities such as

$$\frac{\mu^{\varepsilon/\alpha_1}(dx_1)\mu^{\varepsilon/\alpha_2}(dx_2)}{\mu^{\varepsilon/\alpha_1}(dx_1)\mu^{\varepsilon/\alpha_2}(dx_2) + \mu^{\varepsilon/\alpha_2}(dx_1)\mu^{\varepsilon/\alpha_1}(dx_2)},$$

i.e., likelihood ratios for the corresponding product measure under permutations.

Example 2.1. Consider the case $\mathcal{X} = \mathbb{R}^d$, and

$$\mu^\varepsilon(dx) = \frac{1}{Z^\varepsilon} e^{-\frac{1}{\varepsilon} I(x)} dx, \quad Z^\varepsilon \doteq \int_{\mathbb{R}^d} e^{-\frac{1}{\varepsilon} I(x)} dx, \quad (2.1)$$

where $I : \mathbb{R}^d \rightarrow [0, \infty)$ is a continuous rate function and $\liminf_{M \rightarrow \infty} \inf_{x: \|x\|=M} I(x)/M > 0$ with $\|\cdot\|$ the Euclidean norm (the latter condition in particular guarantees that Z^ε is finite). Then $\{\mu^\varepsilon\}$ automatically satisfies the LDP with rate function I (see Theorem 1.2 in the Appendix). A basic example is $I(x) = \frac{1}{2} \langle x - b, B^{-1}(x - b) \rangle$, where B is a $d \times d$ nonnegative symmetric matrix and $\langle \cdot, \cdot \rangle$ is the standard inner product on $\mathbb{R}^d$, in which case the associated μ^ε is Gaussian with mean b and covariance matrix εB .

Additional examples are given in Section 5. We are interested in estimating *rare event* probabilities such as:

$$\mu^\varepsilon(A) \quad (2.2)$$

for a set $A \subset \mathcal{X}$ that does not contain any of the minimizers of I , and *risk-sensitive* functionals [10] of the form

$$\int_{\mathcal{X}} e^{-\frac{1}{\varepsilon} F^\varepsilon(x)} \mu^\varepsilon(dx) \quad (2.3)$$

for various classes of functions F^ε .

2.2 The INS Estimator

To introduce the estimator, we focus first on the case of risk-sensitive functionals as in (2.3). We also simplify by assuming that $\{\mu^\varepsilon\}$ satisfies conditions similar to Example 2.1, and that $F^\varepsilon = F$ for all ε . Later in Section 4, we outline extensions where the measures do not take this exact form and F^ε depends on ε .

CONDITION 2.2. $\{\mu^\varepsilon\}$ takes the form $\mu^\varepsilon(dx) = \frac{1}{Z^\varepsilon} e^{-\frac{1}{\varepsilon} I(x)} \gamma(dx)$ for a reference measure γ on $\mathcal{X}$ and satisfies the LDP with rate I .

The definition of the INS estimator requires certain parameters, which we now introduce. Given a fixed $K \in \mathbb{N}$, consider numbers $1 = \alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_K > 0$. In a physical system ε would be a parameter that plays the role of temperature, and the parameters $\varepsilon/\alpha_j, j > 1$ define alternative temperatures (larger than the starting temperature since $\alpha_j \leq 1$). The terminology, *infinite swapping*, is inherited from an analogous estimator constructed in a dynamical setting, in which swapping of states allows parallel replicas of the system at different temperatures to be linked. In the present setting, this linkage will be accomplished through appropriate weights. Letting $\alpha \doteq (\alpha_1, \alpha_2, \dots, \alpha_K)$ and $\mathbf{x} \doteq (x_1, x_2, \dots, x_K)$ with each $x_j \in \mathcal{X}$, we define the weights

$$\rho^\varepsilon(\mathbf{x}; \alpha) \doteq \frac{e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_j)}}{\sum_{\tau \in \Sigma_K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_{\tau(j)})}}. \quad (2.4)$$

Let $\{X_j^\varepsilon\}_{j \in \{1, \dots, K\}}$ be independent with $X_j^\varepsilon \sim \mu^{\varepsilon/\alpha_j}$, where $\sim$ means that X_j^ε has the distribution $\mu^{\varepsilon/\alpha_j}$. Let Σ_K denote the set of permutations on $\{1, \dots, K\}$, and for $\sigma \in \Sigma_K$ let $\mathbf{X}_\sigma^\varepsilon$ denote $(X_{\sigma(1)}^\varepsilon, \dots, X_{\sigma(K)}^\varepsilon)$. Then a single sample of the K -temperature INS estimator based on α for estimating $\int_{\mathcal{X}} e^{-\frac{1}{\varepsilon} F(x)} \mu^\varepsilon(dx)$ is

$$\hat{\eta}_{\text{INS}}^\varepsilon(\mathbf{X}^\varepsilon; \alpha) = \sum_{\sigma \in \Sigma_K} \rho^\varepsilon(\mathbf{X}_\sigma^\varepsilon; \alpha) e^{-\frac{1}{\varepsilon} F(X_{\sigma(1)}^\varepsilon)}.$$

Although $\hat{\eta}_{\text{INS}}^\varepsilon$ depends on the number K of temperatures, to simplify notation this is not made explicit. The analogous estimator when applied to the problem of approximating probabilities (i.e., $F(x) = \infty 1_{A^c}(x)$, where 1_{A^c} is the indicator of the complement of A) will be denoted by $\hat{\theta}_{\text{INS}}^\varepsilon$. Since we discuss various unbiased estimators, a subscript (e.g., INS) is sometimes used to distinguish different estimators. If the subscript is missing, then it is always the INS estimator that is meant.

Remark 2.3. We will always assume α satisfies the ordering $1 = \alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_K > 0$.

Before proceeding, we establish basic properties of the estimator.

LEMMA 2.4. Assume Condition 2.2. Then $\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha)$ is an unbiased estimator of $\int_{\mathcal{X}} e^{-\frac{1}{\varepsilon} F(x)} \mu^\varepsilon(dx)$.

PROOF. It follows from the definitions that

$$\begin{aligned} E(\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha)) &= \frac{1}{Z^{\varepsilon; \alpha}} \int_{\mathcal{X}^K} \sum_{\sigma \in \Sigma_K} \rho^\varepsilon(\mathbf{x}_\sigma; \alpha) e^{-\frac{1}{\varepsilon} F(x_{\sigma(1)})} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_j)} \gamma^K(d\mathbf{x}) \\ &= \sum_{\sigma \in \Sigma_K} \frac{1}{Z^{\varepsilon; \alpha}} \int_{\mathcal{X}^K} \frac{e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_{\sigma(j)})} e^{-\frac{1}{\varepsilon} F(x_{\sigma(1)})}}{\sum_{\tau \in \Sigma_K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_{\tau(j)})}} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_j)} \gamma^K(d\mathbf{x}), \end{aligned}$$

where

$$\gamma^K(d\mathbf{x}) \doteq \gamma(dx_1) \cdots \gamma(dx_K) \text{ and } Z^{\varepsilon; \alpha} \doteq Z^\varepsilon Z^{\varepsilon/\alpha_1} \cdots Z^{\varepsilon/\alpha_K}.$$

For each σ , the change of variables from $(x_{\sigma(1)}, \dots, x_{\sigma(K)})$ to $(y_1, \dots, y_K)$ gives:

$$\begin{aligned} & \int_{\mathcal{X}^K} \frac{e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_{\sigma(j)})} e^{-\frac{1}{\varepsilon} F(x_{\sigma(1)})}}{\sum_{\tau \in \Sigma_K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_{\tau(j)})}} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(x_j)} \gamma^K(d\mathbf{x}) \\ &= \int_{\mathcal{X}^K} \frac{e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_j)} e^{-\frac{1}{\varepsilon} F(y_1)}}{\sum_{\tau \in \Sigma_K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_{\tau^{-1}(j)})}} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_{\sigma^{-1}(j)})} \gamma^K(dy). \end{aligned}$$

Thus,

$$\begin{aligned} E(\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha)) &= \sum_{\sigma \in \Sigma_K} \int_{\mathcal{X}^K} \frac{e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_j)} e^{-\frac{1}{\varepsilon} F(y_1)}}{\sum_{\tau \in \Sigma_K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_{\tau^{-1}(j)})}} \frac{e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_{\sigma^{-1}(j)})}}{Z^{\varepsilon; \alpha}} \gamma^K(dy) \\ &= \frac{1}{Z^{\varepsilon; \alpha}} \int_{\mathcal{X}^K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_j)} e^{-\frac{1}{\varepsilon} F(y_1)} \gamma^K(dy) \\ &= \int_{\mathcal{X}} e^{-\frac{1}{\varepsilon} F^\varepsilon(x)} \frac{1}{Z^\varepsilon} e^{-\frac{1}{\varepsilon} I(x)} \gamma(dx). \quad \square \end{aligned}$$

A useful observation is that if $\mathbf{Y}^\varepsilon$ is sampled from the symmetrized version of the distribution of $\mathbf{X}^\varepsilon$, i.e., if

$$\mathbf{Y}^\varepsilon \sim \frac{\sum_{\sigma \in \Sigma_K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_{\sigma(j)})}}{K! Z^{\varepsilon; \alpha}} \gamma^K(dy),$$

then with $\stackrel{d}{=}$ meaning that two random variables have the same distribution, $\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha) \stackrel{d}{=} \hat{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \alpha)$. Indeed, for any $t \in \mathbb{R}$

$$\begin{aligned} E(e^{it\hat{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \alpha)}) &= \int_{\mathcal{X}^K} e^{it\hat{\eta}^\varepsilon(y_1, \dots, y_K; \alpha)} \frac{\sum_{\sigma \in \Sigma_K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_{\sigma(j)})}}{K! Z^{\varepsilon; \alpha}} \gamma^K(dy) \\ &= \frac{1}{K! Z^{\varepsilon; \alpha}} \sum_{\sigma \in \Sigma_K} \int_{\mathcal{X}^K} e^{it\hat{\eta}^\varepsilon(y_1, \dots, y_K; \alpha)} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_j)} \gamma^K(dy) \\ &= \frac{1}{Z^{\varepsilon; \alpha}} \int_{\mathcal{X}^K} e^{it\hat{\eta}^\varepsilon(y_1, \dots, y_K; \alpha)} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I(y_j)} \gamma^K(dy) = E(e^{it\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha)}), \end{aligned}$$

where $i \doteq \sqrt{-1}$ and the fact that $\hat{\eta}^\varepsilon(\mathbf{x}; \alpha)$ is invariant under permutations of $\mathbf{x}$ is used for the second equality.

Remark 2.5. For implementation purposes one prefers $\mathbf{X}^\varepsilon$ over $\mathbf{Y}^\varepsilon$, since it is easier to generate samples of $\mathbf{X}^\varepsilon$ than $\mathbf{Y}^\varepsilon$. However, we introduce $\mathbf{Y}^\varepsilon$ because it is directly analogous to the quantity actually sampled in the dynamical setting which inspired the present static form of INS, and also due to the fact that it will simplify proofs later on. Specifically, we will use that $\{\mathbf{Y}^\varepsilon\}_{\varepsilon \in (0,1)}$ satisfies the LDP with rate function $\min_{\tau \in \Sigma_K} \{\sum_{j=1}^K \alpha_j I(x_{\tau(j)})\}$ (see Lemma 3.8). This rate function gives a simpler expression for the decay rate of $E(\hat{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \alpha))^2$, which in turn allows a simpler proof of our main result concerning the decay rate of $E(\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha))^2$.

3 ANALYSIS OF PERFORMANCE

In this section, we consider the performance of the INS estimator. Since we assume that one can generate independent samples $\{X^\varepsilon(j)\}_{j \in \mathbb{N}}$ from the target measure μ^ε , a possible estimator for (2.2) is the straightforward Monte Carlo (MC) scheme

$$\hat{\theta}_{\text{MC}}^\varepsilon(N) = \frac{1}{N} \sum_{j=1}^N 1_{\{X^\varepsilon(j) \in A\}},$$

where N is the number of samples. The estimator $\hat{\theta}_{MC}^\varepsilon(N)$ is unbiased and by the law of large numbers converges to $\mu^\varepsilon(A)$ almost surely as $N \rightarrow \infty$. However, as is well known, its relative error can make it impractical, since

$$\frac{\sqrt{\text{Var}(\hat{\theta}_{MC}^\varepsilon(N))}}{E[\hat{\theta}_{MC}^\varepsilon(N)]} = \frac{\sqrt{\mu^\varepsilon(A)(1 - \mu^\varepsilon(A))}}{\sqrt{N}\mu^\varepsilon(A)} \approx \frac{1}{\sqrt{N\mu^\varepsilon(A)}}.$$

For small ε , under mild conditions on ∂A (the boundary of A), we have the Laplace type approximation $\mu^\varepsilon(A) \approx e^{-\frac{1}{\varepsilon} \inf_{x \in A} I(x)}$. Thus, to maintain a bounded relative error, one needs exponentially many samples N .

As an alternative, one can consider any estimator of the form

$$\frac{1}{N} \sum_{j=1}^N R_j^\varepsilon, \quad (3.1)$$

with the $\{R_j^\varepsilon\}$ independent and identically distributed (i.i.d.) and $ER_1^\varepsilon = \mu^\varepsilon(A)$ so the estimator is unbiased. Since the variance of (3.1) is proportional to that of R_1^ε , it is enough to consider the single sample estimator R_1^ε when comparing performance. Also, since all estimators under consideration are unbiased, comparing variances reduces to comparing second moments. Similar considerations apply when considering risk-sensitive expected values, for which probabilities can be considered a special case by taking $F(x) = \infty \cdot 1_{A^c}(x)$.

Suppose that $\hat{\eta}^\varepsilon$ is a single sample estimator of $\int_{\mathcal{X}} e^{-\frac{1}{\varepsilon} F(x)} \mu^\varepsilon(dx)$, with F bounded below and continuous. By standard Laplace asymptotics [3, Theorem 1.3.4]

$$I(F) \doteq \inf_{x \in \mathcal{X}} [F(x) + I(x)] = \lim_{\varepsilon \rightarrow 0} -\varepsilon \log \int_{\mathcal{X}} e^{-\frac{1}{\varepsilon} F(x)} \mu^\varepsilon(dx),$$

and the expected value is largely determined by tail properties of μ^ε if $I(x^*) > 0$ for all x^* that minimize in $I(F)$. When this holds, one way of assessing the performance of $\hat{\eta}^\varepsilon$ is to obtain bounds on the limit of $-\varepsilon \log E[(\hat{\eta}^\varepsilon)^2]$ (i.e., the decay rate), or at least lower bounds on

$$\liminf_{\varepsilon \rightarrow 0} -\varepsilon \log E[(\hat{\eta}^\varepsilon)^2],$$

with larger lower bounds indicating better performance. As is well known, there is a limit to how well an estimator can perform. By Jensen's inequality

$$\liminf_{\varepsilon \rightarrow 0} -\varepsilon \log E[(\hat{\eta}^\varepsilon)^2] \leq \liminf_{\varepsilon \rightarrow 0} -2\varepsilon \log \int_{\mathcal{X}} e^{-\frac{1}{\varepsilon} F(x)} \mu^\varepsilon(dx) = 2I(F).$$

An unbiased estimator $\hat{\eta}^\varepsilon$ is said to be *asymptotically optimal* if the upper bound is achieved, i.e.,

$$\lim_{\varepsilon \rightarrow 0} -\varepsilon \log E[(\hat{\eta}^\varepsilon)^2] = 2I(F).$$

As a heuristic motivation for the INS estimator, consider the problem of approximating

$$P(X^\varepsilon \in A) \approx e^{-\frac{1}{\varepsilon} I(a)},$$

where 0 is the unique minimizer of I and $0 \notin A = [a, \infty)$, a is the minimizer of I over $[a, \infty)$, and where $\approx$ indicates that the left- and right-hand sides have the same rate of decay as $\varepsilon \rightarrow 0$. Consider the case of two temperatures. With $\alpha_1 = 1$ and $\alpha_2 = \alpha$

$$\hat{\theta}_{INS}^\varepsilon = \rho^\varepsilon(X_1^\varepsilon, X_2^\varepsilon) 1_A(X_1^\varepsilon) + \rho^\varepsilon(X_2^\varepsilon, X_1^\varepsilon) 1_A(X_2^\varepsilon),$$

where

$$X_1^\varepsilon \sim \frac{1}{Z^\varepsilon} e^{-\frac{1}{\varepsilon} I(x)} dx \text{ and } X_2^\varepsilon \sim \frac{1}{Z^{\varepsilon/\alpha}} e^{-\frac{\alpha}{\varepsilon} I(x)} dx$$

and

$$\rho^\varepsilon(x, y) = \frac{e^{-\frac{1}{\varepsilon}(I(x) + \alpha I(y))}}{e^{-\frac{1}{\varepsilon}(I(x) + \alpha I(y))} + e^{-\frac{1}{\varepsilon}(I(y) + \alpha I(x))}}.$$

Thus, for small $\varepsilon > 0$, there are three types of outcomes:

- $\hat{\theta}_{\text{INS}}^\varepsilon = 1$ when $(X_1^\varepsilon, X_2^\varepsilon) \in A \times A$, which occurs with approximate probability $P((X_1^\varepsilon, X_2^\varepsilon) \in A \times A) \approx e^{-\frac{1}{\varepsilon}I(a)} \cdot e^{-\frac{1}{\varepsilon}\alpha I(a)} = e^{-\frac{1}{\varepsilon}(1+\alpha)I(a)}$.
- $\hat{\theta}_{\text{INS}}^\varepsilon = 0$ when $(X_1^\varepsilon, X_2^\varepsilon) \in A^c \times A^c$, with approximate probability $P((X_1^\varepsilon, X_2^\varepsilon) \in A^c \times A^c) \approx 1$.
- $\hat{\theta}_{\text{INS}}^\varepsilon \approx \rho^\varepsilon(a, 0)$ when $(X_1^\varepsilon, X_2^\varepsilon) \in A \times A^c$ or $A^c \times A$, with approximate probability $P((X_1^\varepsilon, X_2^\varepsilon) \in A \times A^c) + P((X_1^\varepsilon, X_2^\varepsilon) \in A^c \times A) \approx e^{-\frac{1}{\varepsilon}I(a)} + e^{-\frac{1}{\varepsilon}\alpha I(a)} \approx e^{-\frac{1}{\varepsilon}\alpha I(a)}$. (3.2)

Using the definition of ρ^ε gives

$$\rho^\varepsilon(a, 0) = \frac{e^{-\frac{1}{\varepsilon}I(a)}}{e^{-\frac{1}{\varepsilon}I(a)} + e^{-\frac{1}{\varepsilon}\alpha I(a)}} \approx e^{-\frac{1}{\varepsilon}(1-\alpha)I(a)},$$

and therefore

$$\begin{aligned} E(\hat{\theta}_{\text{INS}}^\varepsilon)^2 &\approx 1^2 \cdot e^{-\frac{1}{\varepsilon}(1+\alpha)I(a)} + (\rho^\varepsilon(a, 0))^2 \cdot e^{-\frac{1}{\varepsilon}\alpha I(a)} \\ &\approx e^{-\frac{1}{\varepsilon}(1+\alpha)I(a)} + e^{-\frac{1}{\varepsilon}(2-\alpha)I(a)} \approx e^{-\frac{1}{\varepsilon}[(1+\alpha) \wedge (2-\alpha)]I(a)}. \end{aligned}$$

Thus,

$$\lim_{\varepsilon \rightarrow 0} -\varepsilon \log E(\hat{\theta}_{\text{INS}}^\varepsilon)^2 = ((1+\alpha) \wedge (2-\alpha))I(a),$$

with the maximal rate of decay of $3I(a)/2$ at $\alpha = 1/2$, which can be compared with the rate of decay $I(a)$ that would be found for standard Monte Carlo.

3.1 Approximating a Risk-Sensitive Functional

In this section, we evaluate the decay rate for the INS estimator and the optimal α for a fixed number of temperatures. The first step is to define the function $V_F(\alpha)$ that characterizes the decay rate for a fixed parameter α and solve for the optimum over α . The proof of the following lemma will be given after stating and proving the asymptotic bounds.

LEMMA 3.1. *Let $F : \mathcal{X} \rightarrow \mathbb{R}$ be continuous and bounded below. Let $I : \mathcal{X} \rightarrow [0, \infty]$ be continuous on its domain of finiteness. For fixed $K \in \mathbb{N}$ and α with $1 = \alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_K > 0$, let $D \doteq \{(I(x), F(x)) : x \in \mathcal{X}\}$ and define*

$$V_F(\alpha) \doteq \inf_{\substack{(I_1, F) \in D \\ \{I : J_j \in [0, I_1] \text{ for } j \geq 2\}}} \left[2 \sum_{j=1}^K \alpha_j I_j + 2F - \min_{\sigma \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I_{\sigma(j)} \right\} \right].$$

Then

$$\inf_{x \in \mathcal{X}^K} \left[2 \sum_{j=1}^K \alpha_j I(x_j) + 2F(x_1) - \min_{\sigma \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I(x_{\sigma(j)}) \right\} \right] = V_F(\alpha).$$

Moreover,

$$\sup_{\alpha} V_F(\alpha) = \inf_{x \in \mathcal{X}} \{(2 - (1/2)^{K-1})I(x) + 2F(x)\},$$

and $\alpha^* = (1, 1/2, 1/4, \dots, (1/2)^{K-1})$ achieves the supremum.

When combined with the following result, Lemma 3.1 shows that $\hat{\eta}^\varepsilon(X^\varepsilon; \alpha)$ reaches optimal decay rate when $\alpha = \alpha^*$.

THEOREM 3.2. *Assume Condition 2.2, and let $F : \mathcal{X} \rightarrow \mathbb{R}$ be continuous and bounded below. Let $I : \mathcal{X} \rightarrow [0, \infty]$ be continuous on its domain of finiteness. Then, for any $K \in \mathbb{N}$ and α with $1 = \alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_K > 0$,*

$$\lim_{\varepsilon \rightarrow 0} -\varepsilon \log E(\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha))^2 = V_F(\alpha).$$

Remark 3.3. By formally applying this result with $F(x) = -\infty 1_{A^c}(x)$, we obtain the optimal decay rate for estimating the probability of the set A as in (2.2), a result that will be stated precisely later on. In this case, the decay rate with the optimal α is $(2 - (1/2)^{K-1}) \inf_{x \in A} I(x)$, while the best possible is $2 \inf_{x \in A} I(x)$. With a continuous, nonnegative F , one does even better in some sense, since

$$\inf_{x \in \mathcal{X}} \{(2 - (1/2)^{K-1})I(x) + 2F(x)\} \geq (2 - (1/2)^{K-1})I(F).$$

Remark 3.4. To evaluate a K -temperature INS estimator $\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha)$, we need to sum over $K!$ terms, which, in general, will be computationally expensive. However, for nonnegative F , if $K = 5$ and $\alpha = \alpha^*$, then the decay rate of $E(\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \alpha))^2$ is guaranteed to be at least $1.9375 \cdot I(F)$, while $K = 6$ pushes the rate up to nearly $1.97 \cdot I(F)$. These are both close to the optimum of $2 \cdot I(F)$, and require summing over 120 and 720 terms, respectively. In general, one should balance ε against K to decide on an appropriate value of K .

The proof of Theorem 3.2 will use the next three results. The first extends the standard Laplace principle lower bound and upper bound [3, Section 1.2 and Theorem 1.3.6.], and its proof is in the Appendix.

LEMMA 3.5. *Given a metric space $\mathcal{X}$, suppose $\{X^\varepsilon\} \subset \mathcal{X}$ satisfies the LDP with rate function $I : \mathcal{X} \rightarrow [0, \infty]$. Then, the following conclusions hold:*

- (1) *For any lower semi-continuous function $f : \mathcal{X} \rightarrow \mathbb{R}$ that is bounded below and any closed set $B \subset \mathcal{X}$,*

$$\liminf_{\varepsilon \rightarrow 0} -\varepsilon \log E\left(e^{-\frac{1}{\varepsilon} f(X^\varepsilon)} 1_B(X^\varepsilon)\right) \geq \inf_{x \in B} [f(x) + I(x)].$$

- (2) *Suppose that $\{f_\ell\}_{\ell=1, \dots, N}$ are continuous and $\{g_\ell\}_{\ell=1, 2}$ are lower semi-continuous. Suppose also that $g_r(x) - \min_{\ell \in \{1, \dots, N\}} \{f_\ell(x)\}$ is bounded below for $r = 1, 2$ and that $B_1, B_2 \subset \mathcal{X}$ are closed. Then*

$$\begin{aligned} & \liminf_{\varepsilon \rightarrow 0} -\varepsilon \log E\left(\frac{e^{-\frac{1}{\varepsilon}(g_1(X^\varepsilon) + g_2(X^\varepsilon))}}{\left(\sum_{\ell=1}^N e^{-\frac{1}{\varepsilon} f_\ell(X^\varepsilon)}\right)^2} 1_{B_1}(X^\varepsilon) 1_{B_2}(X^\varepsilon)\right) \\ & \geq \min_{r \in \{1, 2\}} \left\{ \inf_{x \in B_r} \left[2g_r(x) + I(x) - 2 \min_{\ell \in \{1, \dots, N\}} \{f_\ell(x)\} \right] \right\}. \end{aligned}$$

- (3) *Suppose that $I : \mathcal{X} \rightarrow [0, \infty]$ is continuous on its domain of finiteness. Then, for any continuous function $f : \mathcal{X} \rightarrow \mathbb{R}$ that is bounded below and any closed set $B \subset \mathcal{X}$ with B , the closure of its interior,*

$$\limsup_{\varepsilon \rightarrow 0} -\varepsilon \log E\left(e^{-\frac{1}{\varepsilon} f(X^\varepsilon)} 1_B(X^\varepsilon)\right) \leq \inf_{x \in B} [f(x) + I(x)].$$

It will be convenient to exclude the points in the state space where the rate function equals ∞ . The following lemma gives a condition under which this is possible. The proof is in the Appendix.

LEMMA 3.6. *Given a metric space $\mathcal{X}$, suppose that $\{X^\varepsilon\} \subset \mathcal{X}$ satisfies the LDP with rate function $I : \mathcal{X} \rightarrow [0, \infty]$, and assume that*

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon \log P(X^\varepsilon \in \{I = \infty\}) = -\infty.$$

Let $\tilde{X}^\varepsilon \doteq X^\varepsilon 1_{\{I(X^\varepsilon) < \infty\}} + x^* 1_{\{I(X^\varepsilon) = \infty\}}$, where $x^* \in \mathcal{X}$ is any point such that $I(x^*) < \infty$. Then, $\{\tilde{X}^\varepsilon\} \subset \tilde{\mathcal{X}}$ satisfies the LDP with rate function $\tilde{I} : \tilde{\mathcal{X}} \rightarrow [0, \infty)$, where $\tilde{\mathcal{X}} = \mathcal{X} \cap \{I < \infty\}$ and $\tilde{I} = I$ on $\tilde{\mathcal{X}}$.

Remark 3.7. In the case when $X^\varepsilon \sim \mu^\varepsilon$ with $\mu^\varepsilon(dx) = e^{-\frac{1}{\varepsilon}I(x)}\gamma(dx)/Z^\varepsilon$, we find

$$P(X^\varepsilon \in \{I = \infty\}) = \int_{\{I=\infty\}} \frac{1}{Z^\varepsilon} e^{-\frac{1}{\varepsilon}I(x)} \gamma(dx) = \int_{\{I=\infty\}} 0 \gamma(dx) = 0,$$

which gives

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon \log P(X^\varepsilon \in \{I = \infty\}) = -\infty.$$

It follows that, under Condition 2.2, we can assume without loss that $I(x) < \infty$ for each $x \in \mathcal{X}$ and I is continuous on $\mathcal{X}$.

LEMMA 3.8. Suppose $\{X^\varepsilon\}_\varepsilon \subset \mathcal{X}^K$ satisfies the LDP with rate function I , and let Y^ε be a symmetrized version of X^ε , that is:

$$P(Y^\varepsilon \in A) = \frac{1}{K!} \sum_{\sigma \in \Sigma_K} P(X^\varepsilon_{\sigma} \in A) \text{ for all } A \in \mathcal{B}(\mathcal{X}^K),$$

where $\mathcal{B}(\mathcal{X}^K)$ is the Borel σ -algebra of $\mathcal{X}^K$ and

$$X^\varepsilon_{\sigma} \doteq (X^\varepsilon_{\sigma(1)}, \dots, X^\varepsilon_{\sigma(K)}).$$

Then $\{Y^\varepsilon\}_\varepsilon \subset \mathcal{X}^K$ satisfies the LDP with rate function

$$J(\mathbf{x}) = \min_{\sigma \in \Sigma_K} I(x_{\sigma(1)}, \dots, x_{\sigma(K)}).$$

PROOF. Since $\{X^\varepsilon\}_\varepsilon$ satisfies the LDP with rate function I , the Laplace principle holds [3, Section 1.2], and therefore for any bounded and continuous function $g : \mathcal{X}^K \rightarrow \mathbb{R}$

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log E e^{-\frac{1}{\varepsilon}g(X^\varepsilon)} = - \inf_{\mathbf{x} \in \mathcal{X}^K} [g(\mathbf{x}) + I(\mathbf{x})].$$

For such g and any $\sigma \in \Sigma_K$, define $g_\sigma : \mathcal{X}^K \rightarrow \mathbb{R}$ by

$$g_\sigma(x_1, \dots, x_K) = g(x_{\sigma(1)}, \dots, x_{\sigma(K)}),$$

so that g_σ is also bounded and continuous. Then, $g(X^\varepsilon_{\sigma}) = g_\sigma(X^\varepsilon)$, and therefore:

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log E e^{-\frac{1}{\varepsilon}g(X^\varepsilon_{\sigma})} = \lim_{\varepsilon \rightarrow 0} \varepsilon \log E e^{-\frac{1}{\varepsilon}g_\sigma(X^\varepsilon)} = - \inf_{\mathbf{x} \in \mathcal{X}^K} [g_\sigma(\mathbf{x}) + I(\mathbf{x})].$$

Since Y^ε is a symmetrized version of X^ε

$$\varepsilon \log E e^{-\frac{1}{\varepsilon}g(Y^\varepsilon)} = \varepsilon \log \left[\frac{1}{K!} \sum_{\sigma \in \Sigma_K} E e^{-\frac{1}{\varepsilon}g(X^\varepsilon_{\sigma})} \right],$$

and thus with $I_{\sigma^{-1}}(\mathbf{x}) = I(\mathbf{x}_{\sigma^{-1}})$:

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \varepsilon \log E e^{-\frac{1}{\varepsilon}g(Y^\varepsilon)} &= - \min_{\sigma \in \Sigma_K} \left\{ \inf_{\mathbf{x} \in \mathcal{X}^K} [g_\sigma(\mathbf{x}) + I(\mathbf{x})] \right\} \\ &= - \min_{\sigma \in \Sigma_K} \left\{ \inf_{\mathbf{x} \in \mathcal{X}^K} [g(\mathbf{x}) + I_{\sigma^{-1}}(\mathbf{x})] \right\} \\ &= - \inf_{\mathbf{x} \in \mathcal{X}^K} \left[g(\mathbf{x}) + \min_{\sigma \in \Sigma_K} I_{\sigma^{-1}}(\mathbf{x}) \right] \\ &= - \inf_{\mathbf{x} \in \mathcal{X}^K} [g(\mathbf{x}) + J(\mathbf{x})]. \end{aligned}$$

□

PROOF OF THEOREM 3.2. Recall that $\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \stackrel{d}{=} \hat{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \boldsymbol{\alpha})$, and in particular

$$E(\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}))^2 = E(\hat{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \boldsymbol{\alpha}))^2,$$

so the decay rate of $E(\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}))^2$ is the same as that of $E(\hat{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \boldsymbol{\alpha}))^2$. By definition,

$$\begin{aligned} E(\hat{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \boldsymbol{\alpha}))^2 &= E\left(\sum_{\sigma \in \Sigma_K} \rho^\varepsilon(\mathbf{Y}_\sigma^\varepsilon; \boldsymbol{\alpha}) e^{-\frac{1}{\varepsilon} F(\mathbf{Y}_{\sigma(1)}^\varepsilon)}\right)^2 \\ &= \sum_{\sigma \in \Sigma_K} \sum_{\bar{\sigma} \in \Sigma_K} E\left(\rho^\varepsilon(\mathbf{Y}_\sigma^\varepsilon; \boldsymbol{\alpha}) \rho^\varepsilon(\mathbf{Y}_{\bar{\sigma}}^\varepsilon; \boldsymbol{\alpha}) e^{-\frac{1}{\varepsilon} F(\mathbf{Y}_{\sigma(1)}^\varepsilon)} e^{-\frac{1}{\varepsilon} F(\mathbf{Y}_{\bar{\sigma}(1)}^\varepsilon)}\right). \end{aligned}$$

According to Remark 3.7, we may assume that I is finite and continuous.

For $\tau \in \Sigma_K$, define $f_\tau(\mathbf{x}) \doteq \sum_{j=1}^K \alpha_j I(x_{\tau(j)})$, and for $\sigma, \bar{\sigma} \in \Sigma_K$ let

$$g_1(\mathbf{x}) \doteq f_\sigma(\mathbf{x}) + F(x_{\sigma(1)}) \text{ and } g_2(\mathbf{x}) \doteq f_{\bar{\sigma}}(\mathbf{x}) + F(x_{\bar{\sigma}(1)}).$$

Since I and F are finite, continuous and bounded below, g_r and f_τ are also finite, continuous and $g_r(x) - \min_{\ell \in \{1, \dots, N\}} \{f_\ell(x)\}$ is bounded below for $r = 1, 2$.

Applying Lemma 3.5 with $B_1, B_2 = \mathcal{X}^K$, Lemma 3.8, and using the definition of ρ^ε in (2.4) gives

$$\begin{aligned} \liminf_{\varepsilon \rightarrow 0} -\varepsilon \log E\left(\rho^\varepsilon(\mathbf{Y}_\sigma^\varepsilon; \boldsymbol{\alpha}) \rho^\varepsilon(\mathbf{Y}_{\bar{\sigma}}^\varepsilon; \boldsymbol{\alpha}) e^{-\frac{1}{\varepsilon} F(\mathbf{Y}_{\sigma(1)}^\varepsilon)} e^{-\frac{1}{\varepsilon} F(\mathbf{Y}_{\bar{\sigma}(1)}^\varepsilon)}\right) \\ \geq \inf_{\mathbf{x} \in \mathcal{X}^K} \left[2f_\sigma(\mathbf{x}) + 2F(x_{\sigma(1)}) - \min_{\tau \in \Sigma_K} \{f_\tau(\mathbf{x})\}\right] \wedge \inf_{\mathbf{x} \in \mathcal{X}^K} \left[2f_{\bar{\sigma}}(\mathbf{x}) + 2F(x_{\bar{\sigma}(1)}) - \min_{\tau \in \Sigma_K} \{f_\tau(\mathbf{x})\}\right]. \end{aligned}$$

Let $\iota \in \Sigma_K$ denote the identity permutation. Since for any nonnegative sequences $\{a_\varepsilon\}$ and $\{b_\varepsilon\}$

$$\liminf_{\varepsilon \rightarrow 0} -\varepsilon \log(a_\varepsilon + b_\varepsilon) = \min \left\{ \liminf_{\varepsilon \rightarrow 0} -\varepsilon \log a_\varepsilon, \liminf_{\varepsilon \rightarrow 0} -\varepsilon \log b_\varepsilon \right\},$$

we find that

$$\begin{aligned} \liminf_{\varepsilon \rightarrow 0} -\varepsilon \log E(\hat{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \boldsymbol{\alpha}))^2 &\geq \min_{\sigma \in \Sigma_K} \left\{ \inf_{\mathbf{x} \in \mathcal{X}^K} \left[2f_\sigma(\mathbf{x}) + 2F(x_{\sigma(1)}) - \min_{\tau \in \Sigma_K} \{f_\tau(\mathbf{x})\}\right] \right\} \\ &= \inf_{\mathbf{x} \in \mathcal{X}^K} \left[2f_\iota(\mathbf{x}) + 2F(x_\iota) - \min_{\tau \in \Sigma_K} \{f_\tau(\mathbf{x})\}\right], \end{aligned}$$

where the equality holds because every infimum takes the same value.

For the other direction, observe that

$$\rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) = \frac{e^{-\frac{1}{\varepsilon} f_\iota(\mathbf{x})}}{\sum_{\tau \in \Sigma_K} e^{-\frac{1}{\varepsilon} f_\tau(\mathbf{x})}} \geq \frac{e^{-\frac{1}{\varepsilon} f_\iota(\mathbf{x})}}{K! \max_{\tau \in \Sigma_K} (e^{-\frac{1}{\varepsilon} f_\tau(\mathbf{x})})} = \frac{1}{K!} e^{-\frac{1}{\varepsilon} [f_\iota(\mathbf{x}) - \min_{\tau \in \Sigma_K} (f_\tau(\mathbf{x}))]}.$$

For any $\sigma, \bar{\sigma} \in \Sigma_K$, the function

$$f_\sigma(\mathbf{x}) + f_{\bar{\sigma}}(\mathbf{x}) + F(x_{\sigma(1)}) + F(x_{\bar{\sigma}(1)}) - 2 \min_{\tau \in \Sigma_K} (f_\tau(\mathbf{x})),$$

is finite, continuous and bounded below. Thus, applying Part 3 of Lemma 3.5 with $B = \mathcal{X}^K$ gives

$$\begin{aligned} & \limsup_{\varepsilon \rightarrow 0} -\varepsilon \log E \left[\rho^\varepsilon(Y_\sigma^\varepsilon; \alpha) \rho^\varepsilon(Y_{\bar{\sigma}}^\varepsilon; \alpha) e^{-\frac{1}{\varepsilon} F(Y_{\sigma(1)}^\varepsilon)} e^{-\frac{1}{\varepsilon} F(Y_{\bar{\sigma}(1)}^\varepsilon)} \right] \\ & \leq \limsup_{\varepsilon \rightarrow 0} -\varepsilon \log E \left[\exp \left\{ -\frac{1}{\varepsilon} \left[f_\sigma(Y^\varepsilon) + f_{\bar{\sigma}}(Y^\varepsilon) + F(Y_{\sigma(1)}^\varepsilon) + F(Y_{\bar{\sigma}(1)}^\varepsilon) - 2 \min_{\tau} (f_\tau(Y^\varepsilon)) \right] \right\} \right] \\ & \leq \inf_{\mathbf{x} \in \mathcal{X}^K} \left[f_\sigma(\mathbf{x}) + f_{\bar{\sigma}}(\mathbf{x}) + F(x_{\sigma(1)}) + F(x_{\bar{\sigma}(1)}) - \min_{\tau \in \Sigma_K} (f_\tau(\mathbf{x})) \right]. \end{aligned}$$

Since for any nonnegative sequences $\{a_\varepsilon\}$ and $\{b_\varepsilon\}$

$$\limsup_{\varepsilon \rightarrow 0} -\varepsilon \log(a_\varepsilon + b_\varepsilon) \leq \min \left\{ \limsup_{\varepsilon \rightarrow 0} -\varepsilon \log a_\varepsilon, \limsup_{\varepsilon \rightarrow 0} -\varepsilon \log b_\varepsilon \right\},$$

we find

$$\begin{aligned} & \limsup_{\varepsilon \rightarrow 0} -\varepsilon \log E (\hat{\eta}^\varepsilon(Y^\varepsilon; \alpha))^2 \\ & \leq \min_{\sigma, \bar{\sigma} \in \Sigma_K} \left\{ \limsup_{\varepsilon \rightarrow 0} -\varepsilon \log E \left[\rho^\varepsilon(Y_\sigma^\varepsilon; \alpha) \rho^\varepsilon(Y_{\bar{\sigma}}^\varepsilon; \alpha) e^{-\frac{1}{\varepsilon} F(Y_{\sigma(1)}^\varepsilon)} e^{-\frac{1}{\varepsilon} F(Y_{\bar{\sigma}(1)}^\varepsilon)} \right] \right\} \\ & \leq \min_{\sigma, \bar{\sigma} \in \Sigma_K} \left\{ \inf_{\mathbf{x} \in \mathcal{X}^K} \left[f_\sigma(\mathbf{x}) + f_{\bar{\sigma}}(\mathbf{x}) + F(x_{\sigma(1)}) + F(x_{\bar{\sigma}(1)}) - \min_{\tau \in \Sigma_K} (f_\tau(\mathbf{x})) \right] \right\} \\ & = \min_{\sigma \in \Sigma_K} \left\{ \inf_{\mathbf{x} \in \mathcal{X}^K} \left[2f_\sigma(\mathbf{x}) + 2F(x_{\sigma(1)}) - \min_{\tau \in \Sigma_K} (f_\tau(\mathbf{x})) \right] \right\} \\ & = \inf_{\mathbf{x} \in \mathcal{X}^K} \left[2f_i(\mathbf{x}) + 2F(x_1) - \min_{\tau \in \Sigma_K} \{f_\tau(\mathbf{x})\} \right]. \end{aligned}$$

Thus, the upper and lower bounds coincide. We now use Lemma 3.1 to identify the limit as $V_F(\alpha)$, which completes the proof. $\square$

PROOF OF LEMMA 3.1. The first step is to express the space $\mathcal{X}^K$ as $\cup_{\tau \in \Sigma_K} O_\tau$, where

$$O_\tau \doteq \left\{ \mathbf{x} \in \mathcal{X}^K : I(x_{\tau(1)}) \leq I(x_{\tau(2)}) \leq \dots \leq I(x_{\tau(K)}) \right\}.$$

For any $\tau \in \Sigma_K$ there exists $L \in \{1, \dots, K\}$ which depends on τ such that $1 = \tau(L)$. We will use the rearrangement inequality [8, Section 10.2, Theorem 368], which says that if $\mathbf{x} \in O_\tau$, then since α_j is nonincreasing in j the minimum in $\min_{\sigma \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I(x_{\sigma(j)}) \right\}$ is at $\sigma = \tau$. Thus,

$$\begin{aligned} & \inf_{\mathbf{x} \in \mathcal{X}^K} \left[2 \sum_{j=1}^K \alpha_j I(x_j) + 2F(x_1) - \min_{\sigma \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I(x_{\sigma(j)}) \right\} \right] \\ & = \min_{\tau \in \Sigma_K} \left\{ \inf_{\mathbf{x} \in O_\tau} \left[2 \sum_{j=1}^K \alpha_j I(x_j) + 2F(x_1) - \min_{\sigma \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I(x_{\sigma(j)}) \right\} \right] \right\} \\ & = \min_{\tau \in \Sigma_K} \left\{ \inf_{\mathbf{x} \in O_\tau} \left[\sum_{j=1}^K (2\alpha_{\tau(j)} - \alpha_j) I(x_{\tau(j)}) + 2F(x_1) \right] \right\}. \end{aligned}$$

Let $\beta_j \doteq 2\alpha_{\tau(j)} - \alpha_j$, and for each $L \in \{1, \dots, K\}$ define the sets

$$O_\tau^L \doteq \left\{ (x_{\tau(1)}, \dots, x_{\tau(L)}) : \mathbf{x} \in O_\tau \right\}$$

and

$$\bar{O}_\tau^L(\mathbf{y}) \doteq \left\{ (x_{\tau(L)}, \dots, x_{\tau(K)}) : \mathbf{x} \in O_\tau \text{ and } (x_{\tau(1)}, \dots, x_{\tau(L)}) = \mathbf{y} \right\}.$$

Note that for each τ (and using that L is the index such that $\tau(L) = 1$)

$$\begin{aligned} & \inf_{\mathbf{x} \in O_\tau} \left[\sum_{j=1}^K \beta_j I(x_{\tau(j)}) + 2F(x_1) \right] \\ &= \inf_{\mathbf{x} \in O_\tau} \left[\sum_{j=1}^{L-1} \beta_j I(x_{\tau(j)}) + \beta_L I(x_{\tau(L)}) + \sum_{j=L+1}^K \beta_j I(x_{\tau(j)}) + 2F(x_{\tau(L)}) \right] \\ &= \inf_{(y_1, \dots, y_L) \in O_\tau^L} \left[\sum_{j=1}^{L-1} \beta_j I(y_j) + \beta_L I(y_L) + 2F(y_L) + \inf_{(z_L, \dots, z_K) \in \bar{O}_\tau^L(y_1, \dots, y_L)} \left[\sum_{j=L+1}^K \beta_j I(z_j) \right] \right]. \end{aligned}$$

Next, we show that given $(y_1, \dots, y_L)$ (and noting that by definition $z_L = y_L$),

$$\inf_{(z_L, \dots, z_K) \in \bar{O}_\tau^L(y_1, \dots, y_L)} \left[\sum_{j=L+1}^K \beta_j I(z_j) \right] = \left(\sum_{j=L+1}^K \beta_j \right) I(y_L). \quad (3.3)$$

Recall that $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_K > 0$. Therefore, $\beta_K = 2\alpha_{\tau(K)} - \alpha_K \geq 2\alpha_K - \alpha_K = \alpha_K > 0$. More generally, since $\tau(j), \dots, \tau(K)$ are distinct values drawn from $\{1, \dots, K\}$, for each j

$$\beta_j + \dots + \beta_K = 2 \sum_{\ell=j}^K \alpha_{\tau(\ell)} - \sum_{\ell=j}^K \alpha_\ell \geq 2 \sum_{\ell=j}^K \alpha_\ell - \sum_{\ell=j}^K \alpha_\ell > 0.$$

Using $\beta_K > 0$ and the fact that $(z_L, \dots, z_K) \in \bar{O}_\tau^L(y_1, \dots, y_L)$ implies the restriction

$$I(z_L) \leq I(z_{L+1}) \leq \dots \leq I(z_K),$$

we can rewrite the infimum as

$$\begin{aligned} & \inf_{(z_L, \dots, z_K) \in \bar{O}_\tau^L(y_1, \dots, y_L)} \left[\sum_{j=L+1}^K \beta_j I(z_j) \right] \\ &= \inf_{(z_L, \dots, z_K) \in \bar{O}_\tau^L(y_1, \dots, y_L)} \left[\sum_{j=L+1}^{K-2} \beta_j I(z_j) + (\beta_{K-1} + \beta_K) I(z_{K-1}) \right]. \end{aligned}$$

Iterating, we have (3.3). Hence, recalling $D \doteq \{(I(x), F(x)) : x \in \mathcal{X}\}$,

$$\begin{aligned} & \inf_{\mathbf{x} \in \mathcal{X}^K} \left[2 \sum_{j=1}^K \alpha_j I(x_j) + 2F(x_1) - \min_{\sigma \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I(x_{\sigma(j)}) \right\} \right] \\ &= \min_{\tau \in \Sigma_K} \left\{ \inf_{\mathbf{x} \in O_\tau} \left[\sum_{j=1}^K (2\alpha_{\tau(j)} - \alpha_j) I(x_{\tau(j)}) + 2F(x_1) \right] \right\} \\ &= \min_{\tau \in \Sigma_K} \left\{ \inf_{(x_{\tau(1)}, \dots, x_{\tau(L)}) \in O_\tau^L} \left[\sum_{j=1}^{L-1} \beta_j I(x_{\tau(j)}) + \left(\sum_{j=L}^K \beta_j \right) I(x_{\tau(L)}) + 2F(x_{\tau(L)}) \right] \right\} \\ &= \min_{\tau \in \Sigma_K} \left\{ \inf_{\left\{ (I_{\tau(L)}, F), (I_{\tau(L)}, F) \in D \right\}} \left[\sum_{j=1}^{L-1} \beta_j I_{\tau(j)} + \left(\sum_{j=L}^K \beta_j \right) I_{\tau(L)} + 2F \right] \right\}. \end{aligned}$$

The last equality holds because I is continuous on the domain of finiteness.

It remains to show that the last display coincides with $V_F(\alpha)$. Recalling the definition of $V_F(\alpha)$, we have

$$\begin{aligned} V_F(\alpha) &\doteq \inf_{\substack{(I_1, F) \in D \\ \{(I_1, \dots, I_K): I_j \in [0, I_1] \text{ for } j \geq 2\}}} \left[2 \sum_{j=1}^K \alpha_j I_j + 2F - \min_{\sigma \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I_{\sigma(j)} \right\} \right] \\ &= \min_{\tau \in \Sigma_K} \left\{ \inf_{\substack{\{(I_{\tau(L)}, F): (I_{\tau(L)}, F) \in D\} \\ \{(I_{\tau(1)}, \dots, I_{\tau(K)}): I_{\tau(1)} \leq I_{\tau(2)} \leq \dots \leq I_{\tau(K)} \leq I_{\tau(L)}\}}} \left[\sum_{j=1}^K (2\alpha_{\tau(j)} - \alpha_j) I_{\tau(j)} + 2F \right] \right\}. \end{aligned}$$

Since $I \in O_\tau$ implies $I_{\tau(j)} \geq I_{\tau(L)}$ and hence $I_{\tau(j)} = I_{\tau(L)}$ for $L < j \leq K$,

$$\begin{aligned} &\inf_{\substack{\{(I_{\tau(L)}, F): (I_{\tau(L)}, F) \in D\} \\ \{(I_{\tau(1)}, \dots, I_{\tau(K)}): I_{\tau(1)} \leq I_{\tau(2)} \leq \dots \leq I_{\tau(K)} \leq I_{\tau(L)}\}}} \left[\sum_{j=1}^K \beta_j I_{\tau(j)} + 2F \right] \\ &= \inf_{\substack{\{(I_{\tau(L)}, F): (I_{\tau(L)}, F) \in D\} \\ \{(I_{\tau(1)}, \dots, I_{\tau(L-1)}): I_{\tau(1)} \leq I_{\tau(2)} \leq \dots \leq I_{\tau(L)}\}}} \left[\sum_{j=1}^{L-1} \beta_j I_{\tau(j)} + \left(\sum_{j=L}^K \beta_j \right) I_{\tau(L)} + 2F \right]. \end{aligned}$$

This completes the proof of the first part.

To prove the second part, i.e., $\sup_{\alpha} V_F(\alpha) = \inf_x \{(2 - (1/2)^{K-1})I(x) + 2F(x)\}$, first rewrite $V_F(\alpha)$ by noticing that since I_1 is the largest value in the set I ,

$$\min_{\tau \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I_{\tau(j)} \right\}$$

obtains the minimum at some $\tau \in \Sigma_K$ with $\tau(K) = 1$. Therefore

$$V_F(\alpha) = \inf_{\substack{(I_1, F) \in D \\ \{I_j \leq I_1 \text{ for } j \geq 2\}}} \left[(2\alpha_1 - \alpha_K) I_1 + 2 \sum_{j=2}^K \alpha_j I_j + 2F - \min_{\tau \in \Sigma_K, \tau(K)=1} \left\{ \sum_{j=1}^{K-1} \alpha_j I_{\tau(j)} \right\} \right].$$

Suppose we are given any $K - 1$ numbers and assign them to $\{I_j\}_{j=2, \dots, K}$ in a certain order. Then, the value of

$$\min_{\tau \in \Sigma_K, \tau(K)=1} \left\{ \sum_{j=1}^{K-1} \alpha_j I_{\tau(j)} \right\}$$

is independent of the order. But since $\alpha_1 \geq \dots \geq \alpha_K > 0$, by the rearrangement inequality, the smallest value of

$$\sum_{j=2}^K \alpha_j I_j$$

is obtained by taking the I_j , $j \geq 2$ in increasing order. By choosing this ordering of $\{I_j\}_{j=2, \dots, K}$,

$$\min_{\tau \in \Sigma_K, \tau(K)=1} \left\{ \sum_{j=1}^{K-1} \alpha_j I_{\tau(j)} \right\} = \sum_{j=2}^K \alpha_{j-1} I_j.$$

Thus,

$$\begin{aligned} V_F(\boldsymbol{\alpha}) &= \inf_{\substack{(I_1, F) \in D \\ \{1 \leq I_2 \leq \dots \leq I_K \leq I_1\}}} \left[(2\alpha_1 - \alpha_K) I_1 + 2 \sum_{j=2}^K \alpha_j I_j + 2F - \sum_{j=2}^K \alpha_{j-1} I_j \right] \\ &= \inf_{\substack{(I_1, F) \in D \\ \{1 \leq I_2 \leq \dots \leq I_K \leq I_1\}}} \left[(2\alpha_1 - \alpha_K) I_1 + \sum_{j=2}^K (2\alpha_j - \alpha_{j-1}) I_j + 2F \right]. \end{aligned} \quad (3.4)$$

Using summation by parts and $\alpha_1 = 1$, we have

$$V_F(\boldsymbol{\alpha}) = \inf_{\substack{(I_1, F) \in D \\ \{1 \leq I_2 \leq \dots \leq I_K \leq I_1\}}} \left[(2\alpha_1 - \alpha_K) I_1 + \sum_{j=2}^{K-1} \alpha_j (2I_j - I_{j+1}) + 2\alpha_K I_K - I_2 + 2F \right]. \quad (3.5)$$

Since I is a rate function and F is continuous and bounded below, there is $(I_0, F_0) \in D$ such that

$$(2 - (1/2)^{K-1})I_0 + 2F_0 = \inf_x \left[(2 - (1/2)^{K-1})I(x) + 2F(x) \right].$$

Let $\boldsymbol{\alpha}^* \doteq (1, 1/2, \dots, 1/2^{K-1})$ and $\mathbf{I}^* = (I_1^*, \dots, I_K^*)$, with $I_1^* \doteq I_0$, $I_j^* \doteq (1/2)^{K-j+1}I_0$ for $j = 2, \dots, K$. We have the following inequalities, which are explained after the display:

$$\begin{aligned} (2 - (1/2)^{K-1})I_0 + 2F_0 &= \inf_{\substack{(I_1, F) \in D \\ \{1 \leq I_2 \leq \dots \leq I_K \leq I_1\}}} \left[(2\alpha_1^* - \alpha_K^*) I_1 + \sum_{j=2}^K (2\alpha_j^* - \alpha_{j-1}^*) I_j + 2F \right] \\ &= V_F(\boldsymbol{\alpha}^*) \\ &\leq \sup_{\boldsymbol{\alpha}} V_F(\boldsymbol{\alpha}) \\ &\leq \sup_{\boldsymbol{\alpha}} \left[(2\alpha_1 - \alpha_K) I_1^* + \sum_{j=2}^{K-1} \alpha_j (2I_j^* - I_{j+1}^*) + 2\alpha_K I_K^* - I_2^* + 2F_0 \right] \\ &= (2 - (1/2)^{K-1})I_0 + 2F_0. \end{aligned}$$

The first equality follows from $2\alpha_j^* - \alpha_{j-1}^* = 0$ for $j = 2, \dots, K$; the second equality from (3.4); the second inequality is from (3.5); the third equality uses $\alpha_1 = 1$, $2I_j^* - I_{j+1}^* = 0$ for $j = 2, \dots, K$, $-\alpha_K I_1^* + 2\alpha_K I_K^* = 0$ and $I_2^* = (1/2)^{K-1}I_0$. We therefore obtain

$$\sup_{\boldsymbol{\alpha}} V_F(\boldsymbol{\alpha}) = \inf_x \left\{ (2 - (1/2)^{K-1})I(x) + 2F(x) \right\}. \quad \square$$

Remark 3.9. Note that the proof identifies $(\boldsymbol{\alpha}^*, \mathbf{I}^*)$ as a saddle point of a min/max problem.

3.2 Estimating the Probability of Rare Events

We next state the corresponding results for the INS estimator for approximating probabilities as in (2.2). The statements are what one would obtain by letting F approximate $-M1_{A^c}$ and sending $M \rightarrow \infty$. The proofs are similar to but easier than those of the risk-sensitive functional, and thus omitted.

Definition 3.10. Given any $K \in \mathbb{N}$ and $\boldsymbol{\alpha}$, let $\{X_j^\varepsilon\}$ be independent with $X_j^\varepsilon \sim \mu^{\alpha_j/\varepsilon}$ for $j \in \{1, \dots, K\}$. With the weights $\rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha})$ defined according to (2.4), the K -temperature INS estimator for estimating $\mu^\varepsilon(A)$ is

$$\hat{\theta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) = \sum_{\sigma \in \Sigma_K} \rho^\varepsilon(\mathbf{X}_\sigma^\varepsilon; \boldsymbol{\alpha}) 1_A(X_{\sigma(1)}^\varepsilon). \quad (3.6)$$

Following the same argument as for $\hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha})$ but replacing the function $e^{-\frac{1}{\varepsilon}F(x)}$ with $1_A(x)$ gives

- $\hat{\theta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha})$ is an unbiased estimator for $\mu^\varepsilon(A)$,
- $\hat{\theta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \stackrel{d}{=} \hat{\theta}^\varepsilon(\mathbf{Y}^\varepsilon; \boldsymbol{\alpha})$.

THEOREM 3.11. *Let A be the closure of its interior. Suppose that $I : \mathcal{X} \rightarrow [0, \infty]$ is continuous on its domain of finiteness. Then, for any $K \in \mathbb{N}$ and parameter $\boldsymbol{\alpha}$,*

$$\lim_{\varepsilon \rightarrow 0} -\varepsilon \log E \left(\hat{\theta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \right)^2 = V(\boldsymbol{\alpha}) \cdot I(A),$$

where,

$$V(\boldsymbol{\alpha}) = \inf_{\{I: I_1=1, I_j \in [0,1] \text{ for } j=2, \dots, K\}} \left[2 \sum_{j=1}^K \alpha_j I_j - \min_{\sigma \in \Sigma_K} \left\{ \sum_{j=1}^K \alpha_j I_{\sigma(j)} \right\} \right].$$

Moreover, $V(\boldsymbol{\alpha})$ achieves its maximal value of $2 - (1/2)^K$ at $\boldsymbol{\alpha}^* = (1, 1/2, \dots, 1/2^{K-1})$.

4 EXTENSIONS

The theory of Section 3 assumes that $\mu^\varepsilon(dx) = \frac{1}{Z^\varepsilon} e^{-\frac{1}{\varepsilon}I(x)} \gamma(dx)$ and that F does not depend on ε . This covers interesting examples, such as Gibbs measures on a continuous space. Moreover, the basic result can be generalized in various directions, which allows a number of other cases, including measures on discrete spaces, mixtures, and measures on path space. In this section, we discuss some of these generalizations.

Suppose we know only that $\{\mu^\varepsilon\}$ satisfies the LDP with rate function I , and μ^ε has a density function $g^\varepsilon(x)$ for any $\varepsilon > 0$. Define a single sample of the K -temperature INS estimator based on $\boldsymbol{\alpha}$ for estimating $\int_{\mathcal{X}} e^{-\frac{1}{\varepsilon}F^\varepsilon(x)} \mu^\varepsilon(dx)$ to be

$$\tilde{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \doteq \sum_{\sigma \in \Sigma_K} \tilde{\rho}^\varepsilon(\mathbf{X}_\sigma^\varepsilon; \boldsymbol{\alpha}) e^{-\frac{1}{\varepsilon}F^\varepsilon(X_{\sigma(1)}^\varepsilon)},$$

where

$$\tilde{\rho}^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) \doteq \frac{\prod_{j=1}^K g^{\varepsilon/\alpha_j}(x_j)}{\sum_{\tau \in \Sigma_K} \left[\prod_{j=1}^K g^{\varepsilon/\alpha_j}(x_{\tau(j)}) \right]}.$$

As before, one can prove that the estimator is unbiased and if $\mathbf{Y}^\varepsilon$ is a symmetrized version of $\mathbf{X}^\varepsilon$, then

$$\tilde{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \stackrel{d}{=} \tilde{\eta}^\varepsilon(\mathbf{Y}^\varepsilon; \boldsymbol{\alpha}).$$

Moreover, recall that

$$\rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) \doteq \frac{\prod_{j=1}^K e^{-\frac{\alpha_j}{\varepsilon}I(x_j)}}{\sum_{\tau \in \Sigma_K} \left[\prod_{j=1}^K e^{-\frac{\alpha_j}{\varepsilon}I(x_{\tau(j)})} \right]}.$$

THEOREM 4.1. *Let $F : \mathcal{X} \rightarrow \mathbb{R}$ be continuous and bounded below, and let $I : \mathcal{X} \rightarrow [0, \infty]$ be continuous on its domain of finiteness. Suppose that $\varepsilon \log(\tilde{\rho}^\varepsilon/\rho^\varepsilon) \rightarrow 0$ and $F^\varepsilon \rightarrow F$ uniformly on each compact set in $\mathcal{X}^K$ and $\mathcal{X}$, respectively, and that $\inf_{\varepsilon \in (0,1)} \inf_{x \in \mathcal{X}} F^\varepsilon(x) > -\infty$. Then*

$$\lim_{\varepsilon \rightarrow 0} -\varepsilon \log E \left(\tilde{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \right)^2 = V_F(\boldsymbol{\alpha}),$$

where $V_F(\boldsymbol{\alpha})$ is as defined in Lemma 3.1.

PROOF. Using the uniform bounds on F^ε and the F and the fact that I is a rate function, there is a compact set $C \subset \mathcal{X}$ such that

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\tilde{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \right)^2 = \lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\tilde{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2$$

and

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\hat{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \right)^2 = \lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\hat{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2,$$

in the sense that each limit must be equal to the other if one of them exists. We know that

$$\lim_{\varepsilon \rightarrow 0} -\varepsilon \log E \left(\hat{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \right)^2 = V_F(\boldsymbol{\alpha}),$$

and thus it suffices to prove

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\tilde{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2 = \lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\hat{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2.$$

Using the uniform convergence of $\varepsilon \log(\tilde{\rho}^\varepsilon / \rho^\varepsilon)$ and F^ε , for any $\delta > 0$, there is a ζ such that, if $0 < \varepsilon < \zeta$, then

$$\sup_{\mathbf{x} \in C^K} |\varepsilon \log(\tilde{\rho}^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) / \rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}))| < \delta \text{ and } \sup_{x \in C} |F^\varepsilon(x) - F(x)| < \delta.$$

Thus, for any $\mathbf{x} \in C^K$ and $x \in C$:

$$\rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) e^{-\frac{1}{\varepsilon}\delta} < \tilde{\rho}^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) < \rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) e^{\frac{1}{\varepsilon}\delta} \text{ and } F(x) - \delta < F^\varepsilon(x) < F(x) + \delta.$$

Furthermore,

$$\begin{aligned} \tilde{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) &= \sum_{\sigma \in \Sigma_K} \tilde{\rho}^\varepsilon(\mathbf{X}_\sigma^\varepsilon; \boldsymbol{\alpha}) e^{-\frac{1}{\varepsilon} F^\varepsilon(X_{\sigma(1)}^\varepsilon)} \\ &\leq \left(\sum_{\sigma \in \Sigma_K} \rho^\varepsilon(\mathbf{X}_\sigma^\varepsilon; \boldsymbol{\alpha}) e^{-\frac{1}{\varepsilon} F(X_{\sigma(1)}^\varepsilon)} \right) e^{\frac{2}{\varepsilon}\delta} = \hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) e^{\frac{2\delta}{\varepsilon}} \end{aligned}$$

and similarly $\tilde{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) \geq \hat{\eta}^\varepsilon(\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) e^{-\frac{2\delta}{\varepsilon}}$. Putting the estimates together gives

$$\limsup_{\varepsilon \rightarrow 0} -\varepsilon \log E \left(\tilde{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2 \leq \limsup_{\varepsilon \rightarrow 0} -\varepsilon \log E \left(\hat{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2 + 2\delta$$

and

$$\liminf_{\varepsilon \rightarrow 0} -\varepsilon \log E \left(\tilde{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2 \geq \liminf_{\varepsilon \rightarrow 0} -\varepsilon \log E \left(\hat{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2 - 2\delta.$$

Since δ is arbitrary, sending $\delta \rightarrow 0$ shows the existence of

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\tilde{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2$$

and

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\tilde{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2 = \lim_{\varepsilon \rightarrow 0} \varepsilon \log E \left(\hat{\eta}^\varepsilon (\mathbf{X}^\varepsilon; \boldsymbol{\alpha}) 1_{(C)^K} (\mathbf{X}^\varepsilon) \right)^2. \quad \square$$

Remark 4.2. Theorem 4.1 is indeed an extension of Theorem 3.2, since in the setting of Theorem 3.2

$$g^\varepsilon(x) = \frac{1}{Z^\varepsilon} e^{-\frac{1}{\varepsilon} I(x)}, \quad (4.1)$$

and thus for any $\mathbf{x}$, $\tilde{\rho}^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) = \rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha})$. A situation where g^ε not of the form (4.1) arises is with densities that are mixtures. These densities take the form $g^\varepsilon(x) = (\sum_{j=1}^m \omega_j e^{-\frac{1}{\varepsilon} I_j(x)}) / Z^\varepsilon$, where each I_j is a rate function and $\omega_j > 0$ for all j . Since

$$\left(\min_{j \in \{1, \dots, m\}} \omega_j \right) e^{-\frac{1}{\varepsilon} \min_k [I_k(x)]} \leq \sum_{j=1}^m \omega_j e^{-\frac{1}{\varepsilon} I_j(x)} \leq \left(\sum_{j=1}^m \omega_j \right) e^{-\frac{1}{\varepsilon} \min_k [I_k(x)]},$$

this implies

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log(\tilde{\rho}^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) / \rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha})) = 0$$

uniformly in $\mathbf{x}$.

Remark 4.3. By using the allowed ε dependence of F^ε , it is possible to apply the INS estimator to certain classes of discrete random variables. For example, let X have the geometric distribution with parameter p , let $X^\varepsilon = \varepsilon X$, and suppose one wants to estimate $E e^{-\frac{1}{\varepsilon} G(X^\varepsilon)}$ with G continuous and bounded below. We can represent X^ε in the form $X^\varepsilon = f^\varepsilon(Y^\varepsilon)$, where Y^ε is exponentially distributed with parameter λ/ε , $\lambda \doteq -\log(1-p)$ and

$$f^\varepsilon(y) = \varepsilon (\lfloor y/\varepsilon \rfloor + 1),$$

and therefore consider the problem $E e^{-\frac{1}{\varepsilon} F^\varepsilon(Y^\varepsilon)}$, where $F^\varepsilon = G \circ f^\varepsilon$, and both F^ε and Y^ε satisfy the assumptions in Theorem 4.1.

Remark 4.4. Suppose $\{\mu^\varepsilon\}_\varepsilon$ are product measures $\{\mu_1^\varepsilon \times \mu_2^\varepsilon\}_\varepsilon$. Moreover, assume that $\{\mu_1^\varepsilon\}_\varepsilon$ and $\{\mu_2^\varepsilon\}_\varepsilon$ satisfy Condition 2.2 on Polish spaces X_1 and X_2 , with rate functions I_1 and I_2 , respectively. In this case, we can use INS only on subset of variables. Precisely, given any $K \in \mathbb{N}$ and $\boldsymbol{\alpha}$, let $\{X_j^\varepsilon\}$ be independent with $X_j^\varepsilon \doteq (X_{1,j}^\varepsilon, X_{2,j}^\varepsilon) \sim \mu^{\alpha_j/\varepsilon} = \mu_1^{\alpha_j/\varepsilon} \times \mu_2^{\alpha_j/\varepsilon}$ for $j \in \{1, \dots, K\}$. Define

$$\hat{\theta}^\varepsilon(X^\varepsilon; \boldsymbol{\alpha}) = \sum_{\sigma \in \Sigma_K} \rho_1^\varepsilon(X_{1,\sigma}^\varepsilon; \boldsymbol{\alpha}) 1_A(X_{1,\sigma(1)}^\varepsilon, X_{2,\sigma(1)}^\varepsilon)$$

with the new weights

$$\rho_1^\varepsilon(\mathbf{x}; \boldsymbol{\alpha}) \doteq \frac{e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I_1(x_j)}}{\sum_{\tau \in \Sigma_K} e^{-\frac{1}{\varepsilon} \sum_{j=1}^K \alpha_j I_1(x_{\tau(j)})}},$$

where $\mathbf{x} \doteq (x_1, x_2, \dots, x_K)$ and with each $x_j \in X_1$. Note that the weights here depend only on I_1 instead of I . One can show that this estimator is also an unbiased estimator of $\mu^\varepsilon(A)$, and its decay rate is between that of Monte Carlo and INS on full set of variables. We will return to this possibility in the last example of the next section.

5 EXAMPLES

In this section, we explore numerical simulations of the INS estimator for the estimation of rare sets A , as specified in Section 3.2. Under the conditions of Theorem 3.11, the K -temperature estimator $\hat{\theta}_{\text{INS}}$ has a decay rate of $(2 - (1/2)^{K-1})I(A)$ when used to estimate $\mu^\varepsilon(A)$ for a suitable set A . One can also consider the preasymptotic (or level ε) decay rate given by $-\varepsilon \log E(\hat{\theta}_{\text{INS}}^\varepsilon)^2$. This quantity can be estimated using independent samples. If $\hat{\theta}_1^\varepsilon, \dots, \hat{\theta}_N^\varepsilon$ are independent samples of an unbiased estimator of $\mu^\varepsilon(A)$, then the *empirical preasymptotic decay rate* is defined as

$$-\varepsilon \log \left(\frac{1}{N} \sum_{j=1}^N (\hat{\theta}_j^\varepsilon)^2 \right). \quad (5.1)$$

Although the preasymptotic decay rate converges to the true decay rate, it is only in practice that we get a sense of how small $\varepsilon > 0$ must be before we see agreement with the limiting decay rate. One expects the large deviation approximation to be useful for larger values of ε for some choices of A and to need smaller ε for others. We briefly revisit the heuristic asymptotic analysis of Section 3 to elucidate this issue.

First, observe that the choice of $\hat{\theta}_1^\varepsilon, \dots, \hat{\theta}_N^\varepsilon$ which maximizes the quantity (5.1) subject to the constraint $(\hat{\theta}_1^\varepsilon + \dots + \hat{\theta}_N^\varepsilon) = N\mu^\varepsilon(A)$ is given by $\hat{\theta}_j = \mu^\varepsilon(A)$. In other words, for a highly accurate estimator to have a near-optimal decay rate it should generate values which are centered around the target value $\mu^\varepsilon(A)$.

Now consider the case of a two-temperature INS estimator with independent samples X_1^ε and X_2^ε , the second denoting the higher temperature; the case of a general number of temperatures is combinatorially more complicated, but the final principle carries over. It was observed in (3.2) that, when $\hat{\theta}_{\text{INS}}^\varepsilon > 0$, we most often have $\hat{\theta}_{\text{INS}}^\varepsilon = \rho^\varepsilon(X_2^\varepsilon, X_1^\varepsilon)$. This weight can be rewritten as

$$\rho^\varepsilon(x_2, x_1) = \left[1 + \exp\left(\frac{1 - \alpha}{\varepsilon}(I(x_2) - I(x_1))\right) \right]^{-1}. \quad (5.2)$$

For the estimator to have a good decay rate, whenever $X_1^\varepsilon \notin A, X_2^\varepsilon \in A$, the weight $\rho^\varepsilon(X_2^\varepsilon, X_1^\varepsilon)$ needs to be exponentially small as $\varepsilon \rightarrow 0$. It can be seen from (5.2) that $\rho^\varepsilon(x_2, x_1) < 1/2$ is equivalent to $I(x_2) > I(x_1)$, which is plausible since we expect the sample in A to have a higher value of I . It is nevertheless possible for small $\varepsilon > 0$ that $X_1^\varepsilon \notin A, X_2^\varepsilon \in A$ while $I(X_2^\varepsilon) < I(X_1^\varepsilon)$. If they occur with high enough probability, such samples will significantly reduce the preasymptotic decay rate.

Observe that when the target set is of the form $A = \{x : I(x) \geq c\}$ for some $c > 0$, such a mismatch of the weights can never occur since $x_1 \notin A, x_2 \in A$ implies $I(x_2) > I(x_1)$. More generally, sets for which $\{x : I(x) \geq I(A)\} \setminus A$ has small probability relative to $\{x : I(x) \geq I(A)\}$ are in some sense “well described” by the rate function, and their preasymptotic decay rate trends to the asymptotic value faster. On the other hand, the probabilities of sets which are not of this specific form will in general have larger probability, and so having a large preasymptotic decay rate might not be as important. We will explore these issues in the examples.

Next, we define some terminology specific to this section so that we may unambiguously discuss implementation issues for the problem of estimating $P(X^\varepsilon \in A)$. The word, *sample*, will refer to a single random variable X^ε , which could itself be multi- or infinite-dimensional. A *trial* will be a single realization of the K -temperature estimator $\hat{\theta}_{\text{INS}}^\varepsilon$, which requires the generation of K samples, one for each temperature. Using the $N \times K$ samples total, we can run N trials, which are then averaged; this constitutes a *simulation*. We define the INS simulation estimate $\hat{p}_A^\varepsilon$ of $P(X^\varepsilon \in A)$ by

$$\hat{p}_A^\varepsilon \doteq \frac{1}{N} \sum_{j=1}^N \hat{\theta}^\varepsilon(X_j^\varepsilon; \alpha).$$

The standard error $\hat{\sigma}_N$ using N trials is the estimated standard deviation divided by the square root of the number of trials,

$$\hat{\sigma}_N \doteq \frac{1}{\sqrt{N}} \sqrt{\frac{1}{N-1} \sum_{j=1}^N \left(\hat{\theta}^\varepsilon(X_j^\varepsilon; \alpha) - \hat{p}_A^\varepsilon \right)^2}.$$

The standard error is used to provide formal confidence intervals,

$$[\hat{p}_A^\varepsilon - 1.96 \hat{\sigma}_N, \hat{p}_A^\varepsilon + 1.96 \hat{\sigma}_N].$$

Table 1. Probabilities and Expected Number of Successes

ε	10^{-2}	5×10^{-3}	2×10^{-3}	1×10^{-3}
p_A^ε	2.90×10^{-2}	2.54×10^{-3}	3.88×10^{-6}	1.27×10^{-10}
$[Np_A^\varepsilon]$	14,500	1,250	2	0

Table 2. Estimates of Probabilities

$K \setminus \varepsilon$	10^{-2}	5×10^{-3}	2×10^{-3}	1×10^{-3}
1	2.89×10^{-2}	2.61×10^{-3}	0	0
2	2.86×10^{-2}	2.53×10^{-3}	4.01×10^{-6}	5.31×10^{-10}
3	2.87×10^{-2}	2.61×10^{-3}	4.04×10^{-6}	1.47×10^{-10}
4	2.90×10^{-2}	2.55×10^{-3}	3.95×10^{-6}	1.35×10^{-10}
5	2.93×10^{-2}	2.64×10^{-3}	3.78×10^{-6}	1.27×10^{-10}

The relative error $\hat{r}_N$ is the ratio of the standard error and the estimate:

$$\hat{r}_N \doteq \frac{\hat{\sigma}_N}{\hat{p}_A^\varepsilon}.$$

Finally, we will report the *normalized empirical decay rate*

$$\hat{\mathcal{D}}_N^\varepsilon \doteq \frac{-\varepsilon \log \left(\frac{1}{N} \sum_{j=1}^N (\hat{\theta}^\varepsilon(X_j^\varepsilon; \alpha)^2) \right)}{-\varepsilon \log \left(\frac{1}{N} \sum_{j=1}^N \hat{\theta}^\varepsilon(X_j^\varepsilon; \alpha) \right)}.$$

All the examples presented below use the asymptotically optimal choice of α identified in Lemma 3.1.

For each example, we report two times, either in seconds (s) or hours (hrs), which sum to the total time required to compute a simulation. The first time is the *time to generate samples*, which is the time it takes to generate the $N \times K$ samples used for a simulation. This part of the computational effort is inherent to the estimation problem, since we have assumed that independent samples are available. It is up to the user to minimize the time to generate them, and we note that little effort has been made on our part to keep this time small. The second is the *time to compute INS*, which involves evaluating the rate function $I(X^\varepsilon)$ on each sample, computing N trials according to (3.6), and averaging them. As we shall see, K can be chosen to ensure an observably improved decay rate while keeping the time to compute the estimator reasonably small. There is a noticeable tradeoff between the value of K and the time to compute the estimator, with the appropriate choice of K best determined on a case-by-case basis. In general, we do not advocate choosing $K > 5$.

Example 5.1. The first example is chosen to demonstrate the theoretical properties of the INS estimator. We consider the case of i.i.d. one-dimensional standard Gaussians, $\xi_j \sim N(0, 1)$. Set

$$X^\varepsilon \doteq \varepsilon \sum_{j=1}^{\lfloor 1/\varepsilon \rfloor} \xi_j.$$

The quantity of interest will be $p_A^\varepsilon \doteq P(X^\varepsilon \in A)$, where A is the non-convex set $A = (-\infty, -0.25] \cup [0.2, \infty)$. Although simple, owing to the non-convexity of A , this problem requires some care in the design of importance sampling schemes for estimating p_A^ε [7]. We expect the INS estimator to have a high preasymptotic decay rate, because $\{x : I(x) \geq I(A)\} = (-\infty, 0.2] \cup [0.2, \infty)$ and thus $x \notin A, y \in A$ implies $I(y) > I(x)$ with high probability.

Table 3. Normalized Empirical Decay Rate and Times to Compute INS

$K \backslash \varepsilon$	10^{-2}	5×10^{-3}	2×10^{-3}	1×10^{-3}	$K \backslash \varepsilon$	10^{-2}	5×10^{-3}	2×10^{-3}	1×10^{-3}
2	1.32	1.30	1.35	1.41	2	<1s	<1s	<1s	<1s
3	1.46	1.47	1.56	1.62	3	1s	1s	1s	1s
4	1.52	1.54	1.65	1.70	4	17s	19s	18s	18s
5	1.56	1.66	1.71	1.79	5	57s	61s	58s	59s

Left: Normalized empirical decay rate for the probabilities estimated in Table 2. No rate is reported for $K = 1$ temperature since it is always equal to 1. Right: The times to compute INS for each of the corresponding simulations. The time to compute samples is less than 0.1 s for each simulation.

Fig. 1. The conditional distribution of $\hat{\theta}_{\text{INS}}^\varepsilon$ given that the k th temperature landed in A for $k = 3, 4, 5$. Here, $\varepsilon = 5 \times 10^{-3}$ and 10^5 samples were used for each of the $K = 5$ temperatures.

In this case, there are standard deterministic methods to calculate the target probabilities up to machine error. The first row in Table 1 gives the target temperature and the second row shows the (rounded) probability of X^ε landing in the target set A . The third row shows the nearest integer amount of expected successful samples for ordinary Monte Carlo $E \sum_{j=1}^N 1_A(X_j^\varepsilon) = NP(X^\varepsilon \in A)$ when $N = 5 \times 10^5$.

Based on the third row of Table 1, for $N = 5 \times 10^5$ we expect that ordinary Monte Carlo (which corresponds to using a single temperature) will provide some estimate for $\varepsilon = 10^{-2}$ and $\varepsilon = 5 \times 10^{-3}$, but not for the other two values of ε . In Table 2, we compare the estimates obtained from simulations with $K = 1, \dots, 5$ temperatures.

For each pair (ε, K) , Table 3 (left) shows the normalized empirical decay rate $\hat{D}_N^\varepsilon$, while Table 3 (right) shows the time to compute INS for each of these simulations. The time to generate all $N \times K$ samples required for a simulation takes a fraction of a second using MATLAB's built-in function `normrnd`. We see that by combining the estimates from different temperatures the INS estimator can approximate rare events using far fewer samples than ordinary Monte Carlo. This is due to the relatively large number of successes from the higher temperatures, which are properly weighted by $\rho^\varepsilon(\mathbf{x}; \boldsymbol{\alpha})$ to produce an unbiased estimate of $P(X^\varepsilon \in A)$ for the target temperature ε .

In Figure 1, we show the empirical distribution of $\hat{\theta}_{\text{INS}}^\varepsilon$ with $K = 5$ temperatures conditioned on one of the three highest temperatures landing in the target set; the two lower temperatures have few successes and are therefore ignored. We take the negative of the base-10 logarithm for easier readability of the distribution. From the rightmost histogram in Figure 1, one can estimate that the highest temperature landed in A roughly 4×10^4 so at least 40% of the INS samples produced were non-zero. One can also see that the distribution of weights is roughly appropriate, in that the

Fig. 2. Comparison of conditional distributions of INS estimator conditioned on the lowest and highest temperatures landing in the rare event set. In blue with dashed edges is the histogram for the one-dimensional samples, and overlaid in dark gray is the case when 5-dimensional samples are used.

Fig. 3. The empirical decay rate as a function of the inverse temperature. Estimates for $K = 2$ stop at $1/\varepsilon = 1100$ because it becomes computationally unfeasible to generate enough samples of the second temperature landing in the target set.

mode of the histogram times the probability of the k th temperature landing in the target set is on the same order of magnitude as the probability $P(X^\varepsilon \in A)$.

We can further explore this example by taking the Gaussians to be d -dimensional while keeping the target set dependent on only the first coordinate. The rate function becomes the sum of squares $x_k \in \mathbb{R}^d$, $I(x_k) = \sum_{k=1}^d x_{k,j}^2/2$. Using $N = 5 \times 10^5$ total samples spread across $K = 5$ temperatures for $\varepsilon = 5 \times 10^{-3}$, we obtain an estimate of 2.49×10^{-3} , in very close agreement with the theoretical value. Following up on the discussion at the beginning of this section, for this problem the super-level sets of the rate function provide a relatively poor approximation to the target set. As we see in Figure 2, the weights for the low temperatures tend to be too large, and the weights for the high temperatures to be too small. This causes a decrease in the preasymptotic decay rate: the normalized empirical decay rate for the original $d = 1$ was 1.62, while for $d = 5$ it was 1.26. In either case, however, standard Monte Carlo was outperformed with no problem-specific modification of the algorithm.

Finally, since it is relatively cheap to generate independent Gaussian random variables, we can run many simulations for very small $\varepsilon > 0$ to probe the asymptotic behavior of the estimator. In Figure 3, we demonstrate that the normalized preasymptotic decay rates approach $2 - (1/2)^{K-1}$. To avoid the clumping of data points for small $\varepsilon > 0$, we plot the normalized empirical decay rate against the inverse temperature ε^{-1} .

Table 4. Results for Estimating the Probability $\mu^\varepsilon(A_1)$ for Different Values of ε with $N = 10^6$ Trials Using $K = 4$ Temperatures

ε	1/20	1/40	1/80
LDP Prediction	8.49×10^{-2}	7.20×10^{-3}	5.20×10^{-5}
Estimate	5.23×10^{-4}	2.35×10^{-5}	6.14×10^{-8}
Confidence interval	$[5.03, 5.43] \times 10^{-4}$	$[2.08, 2.61] \times 10^{-5}$	$[4.42, 7.86] \times 10^{-8}$
Relative error	1.97×10^{-2}	5.73×10^{-2}	1.43×10^{-1}
Decay rate	1.21	1.24	1.40
Time to generate samples	0.7 hrs	1.3 hrs	2.7 hrs
Time to compute INS	42s	43s	48s

We have $I(A_1) \approx 0.123$.

Table 5. Results for Estimating the Probability $\mu^\varepsilon(A_2)$ for Different Values of ε with $N = 10^6$ Trials Using $K = 4$ Temperatures

ε	1/20	1/40	1/80
LDP Prediction	4.54×10^{-5}	2.06×10^{-9}	4.25×10^{-18}
Estimate	2.54×10^{-4}	1.99×10^{-8}	6.68×10^{-17}
Confidence interval	$[2.51, 2.58] \times 10^{-4}$	$[1.91, 2.08] \times 10^{-8}$	$[5.61, 7.74] \times 10^{-17}$
Relative error	7.39×10^{-3}	2.73×10^{-2}	8.14×10^{-2}
Decay rate	1.51	1.65	1.76
Time to generate samples	0.7 hrs	1.3 hrs	2.7 hrs
Time to compute INS	62s	54s	49s

We have $I(A_2) \approx 0.5$.

Example 5.2. This example is intended to illustrate two points: first, that good performance of the INS estimator is not specific to the one-dimensional family of Gaussians from Example 5.2, and second, that the decay rate can depend significantly on the target set A .

An inherent problem is that, whatever test case we choose, we are largely limited by our assumption that independent samples are available from μ^ε for each $\varepsilon > 0$. Indeed, often the central problem is to generate the samples in the first place. Part of the difficulty is that Markov chain Monte Carlo (MCMC), perhaps the most widely used sampling technique, can be prohibitively slow to converge for sufficiently complicated distributions.

We highlight this issue to ensure there is no confusion as to the use of INS on i.i.d. samples. If one needs to resort to MCMC to generate the samples, then the form of INS studied in this article is not the algorithm of choice. Indeed, the original incarnation of INS studied in [4] was designed as an alternative to an accelerated MCMC method called *replica exchange*. One could view the present case of INS on i.i.d. as an “ideal case” analysis of the original infinite swapping.

Despite its limitations, we will use MCMC to imitate a situation where we might directly have access to independent samples. We subsample the resulting trajectories to obtain samples which are “close enough” to independent for our purposes. We will define a strictly convex rate function $I(x)$ and use the construction in Example 2.1, as this will guarantee the distributions μ^ε defined in (2.1) are unimodal and hence do not severely hinder convergence of MCMC. We take

$$I(x) = (x_1^2 + x_2^2 + x_2 + 1 + x_3^4 + 2x_4^2 + x_4 + 1)^2 - 2.64, \quad (5.3)$$

The global minimum of $I(x)$ occurs at $x_m = (0, -0.5, 0, -0.25)$. This x_m is used to determine the shift of 2.64 to ensure that $I(x_m) = 0$ and $I(x) > 0$ for $x \neq x_m$.

To sample from μ^ε for some $\varepsilon > 0$ using MCMC, we use a Metropolis-Hastings algorithm with proposal distribution

$$q^\varepsilon(x, \cdot) \sim x + \varepsilon Uv,$$

where $U \sim \text{Unif}([-1, 1])$ and v is drawn uniformly from $\{e_1, \dots, e_4\}$, where e_j is the j th basis vector. We sub-sample the simulated trajectory every $8/\varepsilon$ steps, and verify that the histograms for the one-dimensional marginal densities look appropriate. The numerical results are shown in Tables 4 and 5 and are based on $N = 1 \times 10^6$ trials using $K = 4$ temperatures for each trial.

We compare two target sets. Define

$$A_1 = \{x_1 \leq -0.1, x_2 \geq -0.35, x_3 \geq 0, x_4 \geq -0.2\}, \quad A_2 = \{I(x) \geq 0.5\}.$$

The sets are chosen so that the probabilities $\mu^\varepsilon(A_j)$ are on the same order of magnitude for $\varepsilon = 1/20$, even though $I(A_1) < I(A_2)$. We expect the decay rate for estimating the probability of landing in A_1 to be lower than that of A_2 since the set A_1 is an intersection of sets and constitutes only a fraction of the region $\{x : I(x) \geq I(A_1)\}$. In contrast, the set A_2 is explicitly chosen as a super-level set of the rate function, so its weights are more likely to be well-distributed in the sense described in the introduction of this section. In both cases, however, the empirical decay rate is strictly bigger than 1.

Observe that the performances reported in Table 4 are significantly lower than the ones in Table 5. When $\varepsilon = 1/20$, the probabilities $\mu^\varepsilon(A_1)$ and $\mu^\varepsilon(A_2)$ are on the same order of magnitude, though the performance for estimating the latter is markedly higher. This is due to the form of the target set A_2 and is accompanied by a disparity of the probabilities: $\mu^\varepsilon(A_2)$ decreases much faster than $\mu^\varepsilon(A_1)$. Since fewer samples are needed to estimate $\mu^{1/80}(A_1)$ than $\mu^{1/80}(A_2)$, a higher decay rate is of less use for the former, which is reflected in the fact that the relative errors are roughly comparable for the two sets (though slightly smaller for A_1).

Example 5.3. For our final example, we consider the level crossing of a Brownian motion over the finite time interval $[0, 1]$. Let $W(t), t \in [0, 1]$ be a standard Brownian motion with $W(0) = 0$, $EW(t) = 0$ and $EW(t)^2 = t$. Define the scaled process

$$W^\varepsilon(t) = \sqrt{\varepsilon}W(t),$$

and let μ^ε denote the induced measure on the space of trajectories $C([0, 1])$ equipped with the topology of uniform convergence. The sequence of processes $\{W^\varepsilon, \varepsilon > 0\}$ satisfies a large deviations principle on $C([0, 1])$ with rate function

$$I(\phi) = \frac{1}{2} \int_0^1 |\dot{\phi}(s)|^2 ds$$

if $\phi \in C([0, 1])$ is absolutely continuous and $\phi(0) = 0$, and $I(\phi) = +\infty$, otherwise. For some fixed $b > 0$, define the rare event set

$$A_b \doteq \{\phi \in C([0, 1]) : \phi(t) \geq b \text{ for some } t \in [0, 1]\}.$$

One can envision many ways of estimating $\mu^\varepsilon(A_b)$ via simulation. For instance, for $M \geq 1$, let $\mathcal{G}_M : \mathbb{R}^M \rightarrow C([0, 1])$ denote the map

$$(\mathcal{G}_M(z_1, \dots, z_M))(t) = \frac{1}{\sqrt{M}} \left(\sum_{j=1}^{\lfloor Mt \rfloor} z_j + (Mt - \lfloor Mt \rfloor) z_{\lfloor Mt \rfloor + 1} \right). \quad (5.4)$$

The map $\mathcal{G}_M$ defines a linearly interpolated random walk with increments z_j over the interval $[j/M - 1/M, j/M]$ for $j = 1, \dots, M$. If $Z_1, \dots, Z_M$ are i.i.d. normal random variables with zero mean and unit variance, Donsker's invariance principle ensures that $W_M^\varepsilon = \sqrt{\varepsilon} \mathcal{G}_M(Z_1, \dots, Z_M)$ converges in distribution to W^ε as $M \rightarrow \infty$.

The corresponding “discretized” rate function on $C([0, 1])$ for $\{W_M^\varepsilon, \varepsilon > 0\}$ can be given in terms of the rate function $I_M(z_1, \dots, z_M) = \sum_{j=1}^M z_j^2/2$ for $\varepsilon(Z_1, \dots, Z_M)$, and this discretized rate function can in principle be used to compute the weights ρ^ε used in the INS estimator. The problem is that the set A_b is not succinctly described by the increments $(z_1, \dots, z_M)$. Observe that

$$(\sqrt{\varepsilon} \mathcal{G})_M^{-1}(A_b) = \left\{ (z_1, \dots, z_M) \in \mathbb{R}^M : \sum_{j=1}^k z_j \geq \frac{b}{\sqrt{\varepsilon}} \text{ for some } 1 \leq k \leq M \right\}.$$

A standard argument via Jensen’s inequality shows that the trajectory $\phi \in A_b$ which minimizes $I(\phi)$ is a straight line between $(0, 0)$ and $(1, b)$, which means that, with high probability, if the trajectory exceeds b , then all the random variables Z_j should be close to $b/\sqrt{M\varepsilon}$. While $I_M(b/\sqrt{M\varepsilon}, \dots, b/\sqrt{M\varepsilon}) = b^2/2\varepsilon$ is independent of M , $EI_M(Z_1, \dots, Z_M) \approx M/2$. This is consistent with the definition of $I(\phi) = \infty$ for ϕ which are not absolutely continuous, as is the case for paths of the Wiener process (almost surely). This means that, as M increases for fixed $\varepsilon > 0$, there are increasingly many ways for the Z_j ’s to produce a large “cost” (a large I_M value) without landing in $(\sqrt{\varepsilon} \mathcal{G}_M(A_b))^{-1}$. This approximation is not well suited for the INS estimator since all the rate functions have a “weak” dependence on the individual random variables Z_j , with no single one or few largely responsible for determining occurrence of the rare event.

An alternative approximation is Lévy’s construction of Brownian motion using the Schauder functions [9]. Let $\phi_0(t) = t$ and, for $n \geq 1$, let $I(n)$ denote the set of odd integers between 1 and 2^n : $I(1) = \{1\}$, $I(2) = \{1, 3\}$, etc. For $n \geq 1$ and $k \in I(n)$, define

$$\phi_{n,k}(t) = \begin{cases} 2^{\frac{n-1}{2}} \left(t - \frac{k-1}{2^n} \right) & t \in \left[\frac{k-1}{2^n}, \frac{k}{2^n} \right) \\ 2^{\frac{n-1}{2}} \left(\frac{k+1}{2^n} - t \right) & t \in \left[\frac{k}{2^n}, \frac{k+1}{2^n} \right) \\ 0 & \text{else} \end{cases}.$$

Let $\mathcal{T}$ be a space of infinite triangular arrays,

$$\mathcal{T} = \{z = (z_0, z_{1,1}, z_{2,1}, \dots) : z_0, z_{n,k} \in \mathbb{R}, n \geq 1, k \in I(n)\}.$$

For $M \geq 1$, define the map

$$(\mathcal{H}_M(z))(t) = z_0 \phi_0(t) + \sum_{n=1}^M \sum_{k \in I(n)} z_{n,k} \phi_{n,k}(t). \quad (5.5)$$

The preimage of A_b under $\mathcal{H}_M$ can be expressed recursively as follows: Let $\mathcal{Z}_M = (\mathcal{H}_M)^{-1}(A_b)$. Clearly, $\mathcal{Z}_0 = \{z : z_0 \geq b\}$, while for $M \geq 1$,

$$\mathcal{Z}_{M+1} = \mathcal{Z}_M \cup \left(\bigcup_{k \in I(M+1)} \{z : (\mathcal{H}_{M+1}(z))(k/2^{M+1}) \geq b\} \right).$$

This decomposition can be deduced from the piecewise linearity of $\mathcal{H}_M(z)$ and shows that, if $\mathcal{H}_M(z) \in A_b$, then $\mathcal{H}_{M'}(z)$ for all $M' \geq M$. Note that $b\phi_0(t)$ is the most likely escape trajectory, so z_0 is by far the most important contributor to the rare event set. Furthermore, if escape does occur strictly prior to time 1, it is likely to happen close to time 1. For this reason, $z_{n,2^n-1}$ is more important than, say, $z_{n,1}$.

As illustrated with the random walk approximation $\mathcal{G}_M$, we want to avoid including unimportant variables in the rate function. Thus, rather than simulating completely independent trajectories for each temperature, we couple the trajectories and use independent copies of only the first four random variables $Z_0, Z_{1,1}, Z_{2,1}, Z_{2,3}$. As pointed out in Remark 4.4, the estimator remains unbiased but could suffer a decrease in the decay rate. On the other hand, coupling the relatively

unimportant random variables removes a source of variance at the preasymptotic level, and at the same time allows control of discretization error by making M large.

The discretized rate function has

$$I_1(z) = \frac{z_0^2}{2} + \frac{z_{1,1}^2}{2} + \frac{z_{2,1}^2}{2} + \frac{z_{2,3}^2}{2},$$

which is all that is needed to construct the INS estimator; see Remark 4.4. We present the numerical results in Table 6. We first comment on the accuracy of the approximation $\mathcal{H}_M$. For fixed $M \geq 1$, the trajectory $\mathcal{H}_M(z)$ is determined by its values at points $j/2^M$, $j = 0, \dots, 2^M$. Our numerical experiments show that, for fixed $\varepsilon > 0$, the estimate can noticeably vary with M , which is an issue inherent to the approximation of a continuous time process and cannot be resolved by using more temperatures. Nevertheless, the estimator shows reasonable agreement with the LDP prediction. Furthermore, 10^7 standard Monte Carlo runs of the random walk approximation $\mathcal{G}_M$ with $M = 14$ gave estimates of 3.91×10^{-3} and 4.34×10^{-4} for $\varepsilon = 3 \times 10^{-2}$ and 2×10^{-2} , respectively, with no estimate being produced for $\varepsilon = 5 \times 10^{-3}$.

We also remark on the potential for additional computational cost when using Lévy's construction (5.5) of Brownian motion as opposed the usual random walk approximation (5.4). To generate a sample path on $[0, 1]$ with a grid with uniform step size 2^{-M} , both constructions require the generation of $O(2^M)$ i.i.d. one-dimensional Gaussian random variables. From this point of view, the two constructions have comparable computational cost. In practice, Lévy's construction requires the storage and retrieval of the basis functions $\phi_{n,k}$, while the random walk approximation requires only a single call to the cumulative sum function in MATLAB. These implementation issues are relevant, but beyond the scope of our article.

The main message is that the INS estimator can effectively transfer information carried by the higher temperatures to the lower ones even in an infinite-dimensional setting, so long as the target set is represented in an appropriate set of coordinates. An interesting question to explore is whether an appropriate basis can be constructed for general diffusion processes.

6 RELATION OF INS TO OTHER ACCELERATED MC SCHEMES

It is useful to compare the qualitative properties of the INS estimator based on i.i.d. samples with other accelerated Monte Carlo schemes. Schemes that one might consider include standard Monte Carlo, importance sampling, and splitting. Of course, in comparison to the INS approach (at least as developed in this article), these methods apply to a much broader collection of problems. Nonetheless, the comparison is useful in highlighting a few properties of the INS estimator. Our discussion will be brief, and assumes familiarity with importance sampling and standard forms (e.g., fixed rate) of splitting.

Suppose one were to consider, for example, the problem of estimating the probability that a Gaussian random variable X^ε of the form $b + \sqrt{\varepsilon}\sigma Z$ falls into a set A , where b is a d dimensional vector, σ is a $d \times d$ matrix, and Z is a $N(0, I)$ random vector. In order that there be some theoretical bounds on the variance, we assume that the implementation of both importance sampling and splitting are applied to a random walk with each of n summands in the random walk i.i.d. $N(b, \sigma\sigma^T)$ and denoted by $\{Y_j, j = 1, \dots, n\}$, with (for convenience of notation) $\varepsilon = 1/n$ and

$$X_j^n = \frac{1}{n}(Y_1 + \dots + Y_j) \text{ for } j = 1, \dots, n. \quad (6.1)$$

The problem of interest is thus estimating $P(X_n^n \in A)$. We assume the boundary of A is regular enough that $\inf_{x \in A^\circ} I(x) = \inf_{x \in \bar{A}} I(x)$, where A° and $\bar{A}$ denote the interior of A , and the closure of A , respectively. In this case, there are both importance sampling schemes as well as importance functions for the splitting implementation that are known to be asymptotically optimal, so long as

Table 6. Results for Estimating the Probability of Crossing the Threshold $b = 0.5$ for Different Values of ε , with $M = 10$, $N = 5 \times 10^5$ and Using $K = 3$ Temperatures

ε	3×10^{-2}	2×10^{-2}	5×10^{-3}
LDP Prediction	1.55×10^{-2}	1.90×10^{-3}	1.39×10^{-11}
Estimate	3.91×10^{-3}	6.59×10^{-4}	1.45×10^{-11}
Confidence interval	$[3.83, 3.99] \times 10^{-3}$	$[6.23, 6.71] \times 10^{-4}$	$[0.34, 2.57] \times 10^{-11}$
Relative error	1.03×10^{-2}	1.88×10^{-2}	3.92×10^{-1}
Empirical decay rate	1.28	1.29	1.55
Time to generate samples	0.4 hrs	0.4 hrs	0.4 hrs
Time to compute INS	4s	4s	4s

The LDP estimate is $\exp(-b^2/2\varepsilon)$. The computation time reported includes the time taken to generate the samples.

the boundary of A is regular enough [1, 2, 6]. Note that, in contrast with the INS estimator, interpreting X_n^n in terms of the random walk model (6.1) is needed when using these methods to obtain asymptotic optimality. Indeed, the splitting scheme requires that the splitting event be triggered by the entrance of X_j^n into a threshold, whose spacing should be proportional to $1/n$. Likewise, it has been known since [7] that, for the case on non-convex A , one is not able to construct (in general) an asymptotically optimal scheme that considers only a change of measure applied to X^ε (i.e., X_n^n), but rather should consider a dynamic change of measure that depends on how the simulated trajectory evolves [5].

Suppose that we denote estimators obtained using standard Monte Carlo, importance sampling, splitting and the INS schemes by $\hat{\theta}_{MC}^n$, $\hat{\theta}_{IS}^n$, $\hat{\theta}_{Spl}^n$, and $\hat{\theta}_{INS}^n$, respectively. In all cases, we are attempting to approximate a probability θ^n by averaging independent realizations, and each estimator satisfies $E\hat{\theta}_*^n = \theta^n$ and hence is unbiased.

In the case of $\hat{\theta}_{MC}^n$, we attempt to approximate θ^n as a convex combination of 0's and 1's. The variance is approximately the second moment, which decays like the probability itself, i.e., $\exp(-nI(A))$. In this case, the relative error is of course bad as $n \rightarrow \infty$. When $E\hat{\theta}_{MC}^n = P(X_n^n \in A)$ is small it is not the 0's which contribute to the (relatively) large variance, but the occasional 1's. In comparison with standard Monte Carlo, all accelerated schemes attempt to cluster the values taken by $\hat{\theta}_*^n$ (when it is not zero) in a small neighborhood of $P(X_n^n \in A)$.

In the case of importance sampling the estimate takes the form of $R(Y_1, \dots, Y_n)1_{\{X_n^n \in A\}}$, where $R(Y_1, \dots, Y_n)$ is the likelihood ratio of the original distribution with respect to the new distribution used for simulation. For a well-designed scheme, these likelihood ratios will cluster around the target value (approximately $\exp(-nI(A))$), and for this reason, one has good variance reduction and can achieve small relative errors. However, for a poorly designed scheme [7] some samples will produce extraordinarily large likelihood ratios (in fact exponentially large in n), and so it is essential that the design of the simulating distribution be done properly, which in many cases can be challenging. The INS estimator is similar to IS in that it uses weights defined by likelihood ratios in the construction. It differs in that simulations are required for different values of the large deviation parameter, and in that the weights can never be larger than 1.

In the case of splitting, one uses a collection of closed sets $C_0 \subset C_1 \subset \dots \subset C_J = \mathbb{R}^d$ and splitting rates $R_k \in \mathbb{N}$, $k = 0, \dots, J-1$. A single particle is started at some point in $C_J \setminus C_{J-1}$ and evolves according to the law of $\{X_j^n\}$. When a particle enters a set C_k for the first time, it generates $R_k - 1$ successors which evolve independently according to the same law after splitting has occurred. Then, a single sample estimator $\hat{\theta}_{Spl}^n$ is defined to be $N / \prod_{k=0}^{J-1} R_k$, where N in this case is the number of particles simulated which hit A at time n . For a scheme to work well, it

should be designed so that $1/\prod_{k=0}^{J-1} R_k$ is close to the target value, and at the same time one would like at least one particle reach A at time n . To design an effective scheme, one needs to choose the thresholds and splitting rates carefully. In particular, if the thresholds are too far apart, then the entire population will die out in the first few generations with very high probability. On the other hand, if the thresholds are too tight, then the number of particles will grow exponentially which enhance the computational effort exponentially. The fact that both $\hat{\theta}_{\text{Spl}}^n$ and $\hat{\theta}_{\text{INS}}^n$ are confined to $[0, 1]$ means they are in some sense “safer” than IS with regard to variance, though in the case of $\hat{\theta}_{\text{Spl}}^n$, the possibility of exponential growth in the number of particles is in fact as bad as the possibility of exponentially large likelihood ratios.

Thus, INS has features in common with both IS and splitting. However, when applicable, it does not have the same potential for exponentially bad behavior that these schemes can exhibit when poorly designed.

ACKNOWLEDGMENTS

The authors thank an anonymous referee for a number of helpful comments and also suggestions regarding the examples.

REFERENCES

- [1] Jose Blanchet, Peter Glynn, and Kevin Leder. 2012. On Lyapunov inequalities and subsolutions for efficient importance sampling. *ACM Trans. Model. Comput. Simul.* 22, 3, Article 13 (Aug. 2012), 27 pages. DOI : <https://doi.org/10.1145/2331140.2331141>
- [2] T. Dean and P. Dupuis. 2009. Splitting for rare event simulation: A large deviations approach to design and analysis. *Stoch. Proc. Appl.* 119 (2009), 562–587.
- [3] P. Dupuis and R. S. Ellis. 1997. *A Weak Convergence Approach to the Theory of Large Deviations*. John Wiley & Sons, New York.
- [4] P. Dupuis, Y. Liu, N. Plattner, and J. D. Doll. 2012. On the infinite swapping limit for parallel tempering. *SIAM J. Multiscale Model. Simul.* 10 (2012), 986–1022.
- [5] P. Dupuis and H. Wang. 2004. Importance sampling, large deviations, and differential games. *Stoch. and Stoch. Reports.* 76 (2004), 481–508.
- [6] P. Dupuis and H. Wang. 2007. Subsolutions of an Isaacs equation and efficient schemes for importance sampling. *Math. Oper. Res.* 32 (2007), 1–35.
- [7] P. Glasserman and Y. Wang. 1997. Counter examples in importance sampling for large deviations probabilities. *Ann. Appl. Prob.* 7 (1997), 731–746.
- [8] G. H. Hardy, J. E. Littlewood, and G. Pólya. 1952. *Inequalities*. Cambridge Mathematical Library. Cambridge University Press. <https://books.google.com/books?id=t1RCSP8YKt8C>
- [9] I. Karatzas and S. E. Shreve. 1988. *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York.
- [10] P. Whittle. 1981. Risk-sensitive linear/quadratic/Gaussian control. *Adv. Appl. Prob.* 13 (1981), 777–784.

Received November 2017; revised February 2019; accepted March 2019